

IBM DB2 10.1
for Linux, UNIX, and Windows

Installing DB2 Servers

Updated January, 2013



IBM DB2 10.1
for Linux, UNIX, and Windows

Installing DB2 Servers

Updated January, 2013



Note

Before using this information and the product it supports, read the general information under Appendix B, "Notices," on page 561.

Edition Notice

This document contains proprietary information of IBM. It is provided under a license agreement and is protected by copyright law. The information contained in this publication does not include any product warranties, and any statements provided in this manual should not be interpreted as such.

You can order IBM publications online or through your local IBM representative.

- To order publications online, go to the IBM Publications Center at <http://www.ibm.com/shop/publications/order>
- To find your local IBM representative, go to the IBM Directory of Worldwide Contacts at <http://www.ibm.com/planetwide/>

To order DB2 publications from DB2 Marketing and Sales in the United States or Canada, call 1-800-IBM-4YOU (426-4968).

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright IBM Corporation 2013.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Part 1. Installation requirements for DB2 database products 1

Chapter 1. Checking installation prerequisites with db2prereqcheck 3

Part 2. DB2 servers and IBM data server clients 11

Chapter 2. An overview of installing DB2 database servers (Linux and UNIX) 13

Chapter 3. An overview of installing DB2 database servers (Windows) . . . 15

Chapter 4. Non-root installation overview (Linux and UNIX) 17

Installing DB2 database servers as a non-root user	17
Enabling root-based features in non-root installations with db2rfe	18
Differences between root installations and non-root installations	20
Limitations of non-root installations	21
Reducing the size of your DB2 product installation image	23

Chapter 5. Java software support for DB2 database products 27

Chapter 6. Preparing to install DB2 database servers 31

Disk and memory requirements	31
Installation prerequisites for database servers	32
Prerequisites for a DB2 database server installation (Linux and UNIX)	32
Prerequisites for a DB2 database server installation (Windows).	48
Installation requirements for DB2 servers and IBM data server clients (AIX)	56
Installation requirements for DB2 servers and IBM data server clients (HP-UX)	57
Installation requirements for DB2 servers and IBM data server clients (Linux).	61
Installation requirements for DB2 servers and IBM data server clients (Solaris)	66
Installation requirements for DB2 servers and IBM data server clients (Windows)	67
Setting up Network Time Protocol.	70
Configuring hosts as Network Time Protocol clients	70
Setting up a Network Time Protocol server	71

Directory structure for your installed DB2 database product (Windows).	72
Directory structure for your installed DB2 database product (Linux)	77
DB2 Setup wizard language support	81
Supported DB2 interface languages	81
Displaying the DB2 Setup wizard in your national language (Linux and UNIX)	82
Language identifiers for running the DB2 Setup wizard in another language	82
Changing the DB2 interface language (Linux and UNIX)	83
Changing the DB2 product interface language (Windows).	84
Password rules	85
Multiple DB2 copies on the same computer (Windows).	85
Multiple DB2 copies on the same computer (Linux and UNIX)	90
IBM Secure Shell Server For Windows (for Optim Database Administrator usage)	91
Mounting DB2 product media	91
Mounting CDs or DVDs (AIX)	91
Mounting CDs or DVDs (HP-UX)	92
Mounting the CD or DVD (Linux).	93
Mounting CDs or DVDs (Solaris Operating Environment).	93

Chapter 7. Installing DB2 database servers using the DB2 Setup wizard (Windows) 95

Chapter 8. Installing DB2 servers using the DB2 Setup wizard (Linux and UNIX) 99

Chapter 9. DB2 Setup wizard installation walkthrough. 103

Chapter 10. Verifying the installation using the command line processor (CLP) 119

Chapter 11. Taking the first steps after installation 121

Deleting a Firefox browser profile	121
Applying DB2 licenses	121
First Steps interface	122
Post-installation tasks for DB2 database servers (Windows)	123
Adding your user ID to the DB2ADMNS and DB2USERS user groups (Windows)	123
Updating your 32-bit DB2 instances to 64-bit instances (Windows)	123

Validating your DB2 copy	124
Setting up notification and contact lists	124
Changing the default DB2 and default IBM database client interface copy after installation (Windows)	125
IBM data server client connectivity using multiple copies	135
Working with existing DB2 copies	137
Checking for DB2 updates	138
Installing the SSH Server component after installing the DB2 product (Windows)	139
Post-installation tasks for DB2 database servers (Linux and UNIX)	139
Listing DB2 database products installed on your system (Linux and UNIX)	139
Validating your DB2 copy	141
Main menu entries for DB2 tools (Linux)	141

Part 3. Installing the DB2 pureScale Feature 143

Chapter 12. Network topology configuration considerations for DB2 pureScale environments. 145

Chapter 13. Shared storage considerations for DB2 pureScale environments 151

Chapter 14. User-managed file system 155

Chapter 15. DB2 Connect considerations with the DB2 pureScale Feature 157

Chapter 16. DB2 client considerations for the DB2 pureScale Feature 159

Chapter 17. Preparing to install the DB2 pureScale Feature for DB2 Enterprise Server Edition 161

Installing DB2 pureScale Feature (AIX)	161
Installation prerequisites for DB2 pureScale Feature (AIX)	161
Required users for a DB2 pureScale Feature installation	170
Installing and setting up OpenSSH	172
Configuring communication adapter ports and switches (AIX)	173
Configuring PVIDs for a DB2 pureScale instance (AIX)	188
Preinstallation checklist for DB2 pureScale Feature (AIX)	189
Installing DB2 pureScale Feature (Linux)	195
Installation prerequisites for DB2 pureScale Feature (Linux)	195

Required users for a DB2 pureScale Feature installation (Linux)	203
Installing and setting up OpenSSH	204
Configuring communication adapter ports and switches (Linux)	205
Modifying kernel parameters (Linux)	232
Preinstallation checklist for DB2 pureScale Feature (Linux)	234
GDPG infrastructure and GDPG-specific prerequisite conditions	242
DB2 pureScale Feature installation methods	246
Installing the DB2 pureScale Feature using the DB2 Setup wizard (AIX and Linux)	248
Installing a DB2 database product using a response file (Linux and UNIX)	266
Installing DB2 pureScale Feature using the db2_install command	267
Setting up a Network Time Protocol server	269
Configuring hosts as Network Time Protocol clients	270

Chapter 18. Taking the first steps after installing the DB2 pureScale Feature . 273

Verifying your DB2 pureScale installation and setup	273
Manually allowing the DB2 installer to take over a user-managed GPFS cluster.	273
Enabling SCSI-3 PR for DB2 pureScale Feature (Linux)	274
Registering a DB2 product or feature license key using the db2licm command	275
Creating a DB2 pureScale instance	276
Adding communication adapter ports to a CF or member	278
Adding new members or an additional cluster caching facility	280
Updating an instance to a DB2 pureScale instance with the db2iupdt command	281
Applying fix packs	282
Updating an instance to a higher level within a release using the db2iupdt command	285
Enable or disable remote root login	285
Changing the db2sshd user ID to a different user	286
Re-adding a deleted db2sshd user ID	286

Chapter 19. Creating a shared file system 289

Deleting a shared file system	290
Retrieving file system information	291

Chapter 20. Configuring a GDPG environment 293

Getting the cluster installed and running in a GDPG environment	294
Setting up GPFS replication in a GDPG environment.	302
Creating the database in a GDPG environment	308
Configuring the cluster for high availability in a GDPG environment	309

Chapter 21. Removing DB2 data server and the DB2 pureScale Feature 313

Dropping a member or cluster caching facility	313
Manually cleaning a DB2 managed clustered file system	314

Part 4. Installing the DB2 Partitioned Database Environment. 317

Chapter 22. Setting up a partitioned database environment 319

Chapter 23. Preparing to install the DB2 Database Partitioning Feature . . 323

Preparing the environment for a partitioned DB2 server (Windows)	323
Updating environment settings for a partitioned DB2 installation (AIX)	325
Verifying port range availability on participating computers (Linux and UNIX)	327
Creating a DB2 home file system for a partitioned database environment	328
Creating a DB2 home file system for a partitioned database system (AIX)	328
Creating a DB2 home file system for a partitioned database system (HP-UX)	330
Creating a file system for a partitioned database system (Linux)	332
Creating a DB2 home file system for a partitioned database system (Solaris)	333
Verifying that NFS is running (Linux and UNIX)	335
Creating groups and users in a partitioned database environment	336
Creating required users for a DB2 server installation in a partitioned database environment (AIX)	336
Creating required users for a DB2 server installation in a partitioned database environment (HP-UX)	338
Required users for a DB2 pureScale Feature installation (Linux)	339
Creating required users for a DB2 server installation in a partitioned database environment (Solaris Operating System)	340

Chapter 24. Verifying the installation 343

Verifying access to the registry on the instance-owning computer (Windows)	343
Verifying a partitioned database server installation (Linux and UNIX)	343
Verifying a partitioned database environment installation (Windows)	344

Chapter 25. Taking the first steps after the DB2 Database Partitioning Feature installation 347

Setting up a working collective to distribute commands to multiple AIX nodes	347
--	-----

Updating the node configuration file (Linux and UNIX)	348
Format of the DB2 node configuration file.	349
Enabling communications between database partition servers (Linux and UNIX)	355
Enabling the execution of remote commands (Linux and UNIX)	357

Chapter 26. Response files 359

Installing database partition servers on participating computers using a response file (Windows)	359
Installing database partition servers on participating computers using a response file (Linux and UNIX)	360

Part 5. Response file considerations 363

Chapter 27. Response file installation basics 365

Chapter 28. Response file uninstall basics 367

Chapter 29. Response file considerations for a DB2 pureScale environment 369

Chapter 30. Creating a response file using the DB2 Setup wizard or the DB2 Instance Setup wizard 373

Chapter 31. Response file installation of DB2 overview (Linux and UNIX) . . 375

Creating a response file using the sample response file	375
Installing a DB2 database product using a response file (Linux and UNIX)	376
Installing database partition servers on participating computers using a response file (Linux and UNIX)	377
Response file error codes (Linux and UNIX)	378
Uninstalling a DB2 database product, feature, or language using a response file (Linux and UNIX)	379

Chapter 32. Response file installation of DB2 overview (Windows) 381

Making the DB2 installation files available for a response file installation (Windows)	381
Setting up shared access to a directory (Windows)	382
Editing a response file (Windows)	382
Installing a DB2 product using a response file (Windows)	383
Installation of DB2 products using Microsoft Systems Center Configuration Manager (SCCM)	384
Response file error codes (Windows)	384

Response file installation using a batch file (Windows)	385
Uninstalling a DB2 product, feature, or language using a response file (Windows)	386

Chapter 33. The response file generator	387
--	------------

Chapter 34. Available sample response files	389
--	------------

Chapter 35. Sample response files (Linux, UNIX, and Windows)	391
---	------------

Chapter 36. Response file keywords	393
---	------------

Chapter 37. Embedding the DB2 installation image (Linux and UNIX)	411
--	------------

Chapter 38. Exporting and importing a profile	413
--	------------

Chapter 39. Stopping DB2 processes during an interactive installation (Windows)	415
--	------------

Chapter 40. Stopping DB2 processes during a response file installation (Windows)	417
---	------------

Part 6. Installing the DB2 Information Center	419
--	------------

Chapter 41. Installing the DB2 Information Center using the DB2 Setup wizard (Linux)	421
---	------------

Chapter 42. Installing the DB2 Information Center using the DB2 Setup wizard (Windows)	425
---	------------

Chapter 43. DB2 Information Center installation options	427
--	------------

Chapter 44. Starting or stopping the Information Center (Linux and Windows)	429
--	------------

Chapter 45. Manually updating the DB2 Information Center installed on your computer or intranet server	431
---	------------

Part 7. Converting instances to a new DB2 database product	433
---	------------

Chapter 46. Converting instances from one DB2 database product to another (Windows)	435
--	------------

Chapter 47. Converting instances from one DB2 database product to another (Linux and UNIX)	437
---	------------

Part 8. Convert instances to the IBM DB2 pureScale Feature	439
---	------------

Chapter 48. Planning to convert your environment to DB2 pureScale	441
--	------------

Chapter 49. Pre-conversion tasks for DB2 pureScale environments	443
--	------------

Setting up a GPFS file system for a DB2 pureScale environment.	443
Converting table spaces for a DB2 pureScale environment.	445
Verifying that your databases are ready for a DB2 pureScale environment	446
Backing up databases when converting to a DB2 data sharing environment	447

Chapter 50. Converting your existing DB2 instances to a DB2 pureScale environment	449
--	------------

Chapter 51. Post-conversion tasks for a DB2 pureScale environment	451
--	------------

Adjusting database configuration parameters to meet DB2 pureScale environment requirements	451
Backing up databases when converting to a DB2 data sharing environment	452

Part 9. Applying fix packs	455
---	------------

Chapter 52. Preparing to install a fix pack	459
--	------------

Checking fix pack prerequisites	460
Getting fix packs	461
Uncompressing fix packs (Linux and UNIX)	462
Uncompressing fix packs (Windows)	462
Stopping all DB2 processes (Linux and UNIX)	463
Stopping all DB2 instances, services and applications (Windows)	464

Chapter 53. Fix pack installation methods	467
--	------------

Universal versus product-specific fix packs	467
Installing a fix pack (Linux and UNIX)	469

Installing a fix pack to update a Version 10.1 DB2 pureScale instance	469
Installing a fix pack to update existing DB2 database products (Linux and UNIX)	473
Installing a fix pack to install new DB2 database products (Linux and UNIX)	473
Installing a fix pack (Windows)	475
Installing a fix pack for a single database product (Windows)	475
Installing a fix pack for multiple database products (Windows)	475
Installing a fix pack using a response file (Windows)	476
Installing a fix pack in a Microsoft Cluster Server environment (Windows)	477

Chapter 54. After installing a fix pack	481
Post-installation tasks for fix packs (Windows)	481
Post-installation tasks for fix packs (Linux and UNIX)	482
Binding bind files after installing fix packs	484

Chapter 55. Uninstalling fix packs	487
---	------------

Part 10. Uninstalling DB2 database products 489

Chapter 56. Uninstalling your DB2 database product (Linux and UNIX)	491
Stopping the DB2 administration server (Linux and UNIX)	491
Removing the DB2 administration server (Linux and UNIX)	492
Stopping root DB2 instances (Linux and UNIX)	492
Removing DB2 instances (Linux and UNIX)	493
Removing DB2 database products using the db2_deinstall and doce_deinstall commands (Linux and UNIX)	494

Chapter 57. Uninstalling non-root DB2 database products (Linux and UNIX)	497
Stopping non-root instances (Linux and UNIX)	497
Removing non-root DB2 database products using db2_deinstall (Linux and UNIX)	498

Chapter 58. Uninstalling your DB2 database product (Windows).	499
--	------------

Chapter 59. Uninstalling DB2 copies and IBM database client interface copies	501
---	------------

Chapter 60. Uninstalling a DB2 database product, feature, or language using a response file (Linux and UNIX)	509
---	------------

Chapter 61. Uninstalling a DB2 product, feature, or language using a response file (Windows).	511
--	------------

Chapter 62. Uninstalling fix packs	513
---	------------

Part 11. Installing and upgrading SA MP with the DB2 installer 515

Chapter 63. IBM Tivoli System Automation for Multiplatforms (SA MP) base component	517
---	------------

Chapter 64. Installing IBM Tivoli System Automation for Multiplatforms (SA MP).	519
Installing IBM Tivoli System Automation for Multiplatforms (SA MP) using the DB2 installer	519
Installing IBM Tivoli System Automation for Multiplatforms (SA MP) using the installSAM install script	521

Chapter 65. Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP).	523
Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP) using the DB2 installer	524
Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP) using the installSAM installation script	525

Chapter 66. Uninstalling IBM Tivoli System Automation for Multiplatforms (SA MP).	527
Uninstalling IBM Tivoli System Automation for Multiplatforms (SA MP) using the DB2 installer	527
Uninstalling IBM Tivoli System Automation for Multiplatforms (SA MP) using the uninstallSAM uninstall script	528

Chapter 67. Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP).	531
Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP) using DB2 installer	531
Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP) manually	532

Chapter 68. IBM Tivoli System Automation for Multiplatforms (SA MP) install and uninstall logs	535
---	------------

Chapter 69. License terms for using IBM Tivoli System Automation for Multiplatforms (SA MP) integrated with IBM DB2 server	537
---	------------

Chapter 70. Supported software and hardware for IBM Tivoli System Automation for Multiplatforms (SA MP).	539
---	------------

Part 12. IBM Data Studio integrated installation with DB2 database products	541
--	------------

Chapter 71. Installing IBM Data Studio using the DB2 Setup wizard (Windows).	543
---	------------

Chapter 72. Installing IBM Data Studio using the DB2 Setup wizard (Linux).	545
---	------------

Chapter 73. Scenario: Installation of IBM Data Studio client	547
---	------------

Part 13. Appendixes	549
--------------------------------------	------------

Appendix A. Overview of the DB2 technical information	551
--	------------

DB2 technical library in hardcopy or PDF format	551
Displaying SQL state help from the command line processor	554
Accessing different versions of the DB2 Information Center	554
Updating the DB2 Information Center installed on your computer or intranet server	554
Manually updating the DB2 Information Center installed on your computer or intranet server	556
DB2 tutorials	558
DB2 troubleshooting information	558
Terms and conditions.	558

Appendix B. Notices	561
--------------------------------------	------------

Index	565
------------------------	------------

Part 1. Installation requirements for DB2 database products

Before you install your DB2® database product, ensure that the system you choose meets the necessary operating system, hardware, software, communications, disk and memory requirements. The **db2prereqcheck** command checks whether your system meets the prerequisites for the installation.

“Disk and memory requirements” on page 31

The disk space required for your product depends on the type of installation you choose and the type of file system you have. Likewise, memory requirements are affected by the size and complexity of your database system.

DB2 database and IBM data server client products

There is a list of operating system, software, and hardware prerequisites that must be met for each operating system, as follows:

- “Installation requirements for DB2 servers and IBM data server clients (AIX)” on page 56
- “Installation requirements for DB2 servers and IBM data server clients (HP-UX)” on page 57
- “Installation requirements for DB2 servers and IBM data server clients (Linux)” on page 59
- “Installation requirements for DB2 servers and IBM data server clients (Solaris)” on page 66
- “Installation requirements for DB2 servers and IBM data server clients (Windows)” on page 67

Installation requirements for DB2 Connect™ products

There is a list of operating system, software, and hardware prerequisites that must be met for each operating system, as follows:

- “Installation requirements for DB2 Connect Personal Edition (Windows)” in *DB2 Connect Installing and Configuring DB2 Connect Personal Edition*
- “Installation requirements for DB2 Connect Personal Edition (Linux)” in *DB2 Connect Installing and Configuring DB2 Connect Personal Edition*
- “Installation requirements for DB2 Connect products (Solaris Operating System)” in *DB2 Connect Installing and Configuring DB2 Connect Servers*
- “Installation requirements for DB2 Connect server products (Windows)” in *DB2 Connect Installing and Configuring DB2 Connect Servers*
- “Installation requirements for DB2 Connect server products (Linux)” in *DB2 Connect Installing and Configuring DB2 Connect Servers*
- “Installation requirements for DB2 Connect server products (AIX®)” in *DB2 Connect Installing and Configuring DB2 Connect Servers*
- “Installation requirements for DB2 Connect server products (HP-UX)” in *DB2 Connect Installing and Configuring DB2 Connect Servers*

IBM® DB2 pureScale® Feature

There is a list of operating system, software, hardware, and firmware prerequisites that must be met for each operating system, as follows:

- “Installation prerequisites for DB2 pureScale Feature (AIX)” on page 161
- “Installation prerequisites for DB2 pureScale Feature (Linux)” on page 195

Chapter 5, “Java software support for DB2 database products,” on page 27

You require the appropriate level of IBM Software Development Kit (SDK) for Java™ to use Java-based tools and to create and run Java applications, including stored procedures and user-defined functions.

Virtualization support

For details about virtualization technology supported by DB2 database products, see <http://www.ibm.com/developerworks/wikis/display/im/DB2+Virtualization+Support>.

Installation requirements for the DB2 National Language Pack (NLPACK)

On Linux and UNIX operating systems, before installing the DB2 NLPACK, a DB2 server or client product must already be installed. This list of qualified servers or clients include:

- DB2 Enterprise Server Edition
- DB2 Advanced Enterprise Server Edition
- DB2 Workgroup Server Edition
- DB2 Express® Edition
- DB2 Express-C
- DB2 Connect Enterprise Edition
- DB2 Connect Personal Edition
- Data Server Client
- Data Server Runtime Client

Installation requirements for DB2 Spatial Extender

- “System requirements for installing Spatial Extender” in *Spatial Extender User's Guide and Reference*

Installation requirements for InfoSphere® Federation Server products

- Hardware and software requirements for Federation Server at <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.swg.im.iis.prod.install.core.doc/topics/iypisrq-sys.html>
- Hardware and software requirements for Replication Server or Data Event Publisher at <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/topic/com.ibm.swg.im.iis.repl.install.doc/topics/iypirrq-sys.html>

Chapter 1. Checking installation prerequisites with db2prereqcheck

You can use the **db2prereqcheck** command to check the software and firmware prerequisites of a specific DB2 version.

With the **db2prereqcheck** command, you can determine if your system satisfies the DB2 installation prerequisite without having to start the installation of a specific DB2 version. This command checks prerequisites such as operating system level, Linux distribution, AIX Technology Level, C library & runtime patch, Client, uDAPL, GPL or DB2 pureScale specific requirements.

The DB2 product comes with the db2prereqcheck binary and the DB2prereqs.xml file in the cfg folder. The DB2prereqs.xml file contains the prerequisite from Version 9.8 Fix Pack 2 to Version 10.1, and the contents should not be modified. Before you can run the **db2prereqcheck** command, you must extract the binary file and the XML file to your temporary folder in your desktop. You do not need to extract the files if your operating system has DB2 product installed. You need read and write access to the system to run **db2prereqcheck** command (that is, read access for reading in DB2prereqs.xml and write access to generate the output file).

Note: For UNIX operating system, you must extract the db2prereqcheck_exe binary, the DB2prereqs.xml file and the db2prereqcheck shell script.

If you run the **db2prereqcheck** command with no parameters, it determines whether the system meets the prerequisites for all the DB2 releases and fix packs that are defined in the resource XML file.

You can check prerequisites requirement for the latest DB2 version, a specific DB2 version, generate a report file, and more with the **db2prereqcheck** command.

- To check the prerequisites for the latest DB2 version that is defined in the resource XML file, run **db2prereqcheck** command with **-i** parameter. For example,

```
$ ./db2prereqcheck -i
```
- To check whether the system meets the prerequisite for the specific DB2 version that is defined in the resource XML file, run **db2prereqcheck** command with **-v <version number>** parameter. For example,

```
$ ./db2prereqcheck -v 10.1.0.0
```
- To generate a report file containing validation information and output from the **db2prereqcheck** command (including a summary section listing all failing tests), use the **-o <filename>** parameter. Specify the full report filename path to generate the report in another directory. The prerequisites for installing the DB2 product will only be displayed on the command screen without the **-o <filename>** parameter.

```
$ ./db2prereqcheck -i -o report.rpt
```
- The default XML file (with file path : DB2 installation/cfg/DB2prereqs.xml) will be used if the **-f** parameter is not specified. If this default XML file is not found, **-f <filename>** can be used to specify the XML file.

```
$ ./db2prereqcheck -i -f ./newDB2prereqs.xml
```
- On supported AIX and Linux operating systems, to validate prerequisites requirement for a DB2 pureScale installation, use the **-p** option.

```
$ ./db2prereqcheck -i -p
```

- On supported Linux operating systems, to validate prerequisites requirement for a specific network configuration of a DB2 pureScale installation, use the **db2prereqcheck** command with both the **-p** and **-t <network configuration type>** options. The network configuration type should be set to either single InfiniBand port cluster, multiple InfiniBand port cluster, single RoCE port cluster or multiple RoCE port cluster (**SINGLE_IB_PORT_CLUSTER**, **MULT_IB_PORT_CLUSTER**, **SINGLE_ROCE_PORT_CLUSTER** or **MULTI_ROCE_PORT_CLUSTER**).

```
$ ./db2prereqcheck -i -p -t <network configuration type>
```

Note: For a DB2 pureScale installation on AIX using 10GE (RoCE) networks, the AIX and uDAPL software prerequisites must be validated manually. The **db2prereqcheck** command does not automatically validate these levels.

If the **db2prereqcheck** command indicates that the system fails to meet the prerequisite, DBT3507E, for example,

```
Validating ofed ...
Required minimum version and release for ofed: 1.5.2-(null)
DBT3507E The db2prereqcheck utility failed to verify installation prerequisites
because the utility failed to find the following package or file: "ofed".
ERROR: Requirement not matched.
```

check the requirement page under related reference and run it again. For more information, search the DBT message (for example DBT3507E) on Version 10.1 Information Center.

Example

Sample **db2prereqcheck** command output on AIX operating system. For DB2 pureScale installation, use the **-p** option.

```
> db2prereqcheck -p -v 10.1.0.0
=====
Checking DB2 prerequisites for DB2 database version 10.1.0.0 on operating
system "AIX"

Validating kernel level ...
Required minimum operating system kernel level : "6.1".
Actual operating system kernel level: "6.1".
Requirement matched.

Validating AIX technology level and service pack ...
Required minimum technology level: 6 Service pack: 5
Requirement matched.

Validating XL C/C++ Runtime ...
Required minimum XL C/C++ Runtime level: 11.1.0.1
Actual version: 11.1.0.1
Requirement matched.

Validating uDAPL ...
Required minimum uDAPL level: 6.1.6.1
Actual version: 6.1.6.1
Requirement matched.

DBT3533I The db2prereqcheck utility has confirmed that all installation
prerequisites were met for DB2 database server with DB2 pureScale
feature. Version: "10.1.0.0".
```

Sample **db2prereqcheck** command output on AIX operating system. For DB2 pureScale uDAPL only, use the **-u** option.

```

> db2prereqcheck -u -v 10.1.0.0
=====
Checking DB2 prerequisites for DB2 database version 10.1.0.0 on operating
system "AIX"

Validating uDAPL ...
  Required minimum uDAPL level:  6.1.6.1
  Actual version:  6.1.6.1
  Requirement matched.

DBT3533I  The db2prereqcheck utility has confirmed that all installation
prerequisites were met for DB2 database server . Version: "10.1.0.0".

=====
Checking DB2 prerequisites for DB2 database version 10.1.0.0 on operating
system "AIX"

Validating uDAPL ...
  Required minimum uDAPL level:  6.1.6.1
  Actual version:  6.1.6.1
  Requirement matched.

DBT3533I  The db2prereqcheck utility has confirmed that all installation
prerequisites were met for DB2 database server with DB2 pureScale
feature. Version: "10.1.0

```

Sample **db2prereqcheck** command output on AIX operating system. For client installation, use the **-c** option.

```

> db2prereqcheck -c -v 10.1.0.0
=====
Checking DB2 prerequisites for DB2 database version 10.1.0.0 on operating
system "AIX"

Validating kernel level ...
  Required minimum operating system kernel level : "6.1".
  Actual operating system kernel level: "6.1".
  Requirement matched.

Validating AIX technology level and service pack ...
  Required minimum technology level: 6 Service pack: 5
  Requirement matched.

Validating XL C/C++ Runtime ...
  Required minimum XL C/C++ Runtime level:  11.1.0.1
  Actual version:  11.1.0.1
  Requirement matched.

DBT3533I  The db2prereqcheck utility has confirmed that all installation
prerequisites were met for DB2 database client . Version: "10.1.0.0".

=====
Checking DB2 prerequisites for DB2 database version 10.1.0.0 on operating
system "AIX"

Validating kernel level ...
  Required minimum operating system kernel level : "6.1".
  Actual operating system kernel level: "6.1".
  Requirement matched.

Validating AIX technology level and service pack ...
  Required minimum technology level: 6 Service pack: 5
  Requirement matched.

Validating XL C/C++ Runtime ...
  Required minimum XL C/C++ Runtime level:  11.1.0.1
  Actual version:  11.1.0.1

```


Requirement matched.

Validating uDAPL ...

Required minimum uDAPL level: 6.1.6.1

Actual version: 6.1.6.1

Requirement matched.

DBT3533I The db2prereqcheck utility has confirmed that all installation prerequisites were met for DB2 database client with DB2 pureScale feature. Version: "10.1.0.0".

Sample **db2prereqcheck** command output on AIX operating system. To display validation summary only, use the **-s** option.

```
> db2prereqcheck -s -v 10.1.0.0
```

DBT3533I The db2prereqcheck utility has confirmed that all installation prerequisites were met for DB2 database server . Version: "10.1.0.0".

DBT3533I The db2prereqcheck utility has confirmed that all installation prerequisites were met for DB2 database server with DB2 pureScale feature. Version: "10.1.0.0".

Sample **db2prereqcheck** command output on Linux operating system. For DB2 pureScale installation, use the **-p** option.

```
> db2prereqcheck -p -v 10.1.0.0
```

```
=====
Checking DB2 Prerequisites for DB2 Version 10.1.0.0 on Linux
```

Validating Linux distribution ...

Required minimum SUSE version : (null) SP (null)

Validating sles-release ...

Required minimum version and release for sles-release: 10-15.45.8

Actual version of package: 11.1

Requirement matched.

Validating kernel level ...

Required minimum kernel level : 2.6.16

Actual kernel level: 2.6.32.45

Requirement matched.

Validating C++ Library version ...

Required minimum C++ library: libstdc++.so.6

Standard C++ library is located at: /usr/lib64/libstdc++.so.6.0.10

Actual C++ library: CXXABI_1.3.1

Validating 32 bit version of libstdc++.so.6 ...

Found the 64 bit /usr/lib64/libstdc++.so.6 at /usr/lib64

Found the 32 bit /usr/lib/libstdc++.so.6 at /usr/lib

Requirement matched.

Validating libc.so version ...

glibc library is located at: /lib64/libc-2.11.1.so

Required minimum glibc library version: 2.4.0

Actual glibc library version: 2.11.1

Requirement matched.

Validating libaio.so version ...

Loading of libaio.so.1 on system succeed.

Requirement matched.

Validating ofed ...

Required minimum version and release for ofed: 1.5.2-(null)

Requirement matched.

```
Validating sg3_utils ...
Package (or file) found: sg3_utils
Requirement matched.
```

```
Validating sg_persist ...
Package (or file) found: /usr/bin/sg_persist
Requirement matched.
```

```
DBT3533I The db2prereqcheck utility has confirmed that all installation
prerequisites were met for DB2 database server with DB2 pureScale feature.
Version: "10.1.0.0".
```

Sample **db2prereqcheck** command output on Linux operating system. Use **-p** and **-t** option to validate prerequisites requirement for a specific network configuration. For example, to validate for a single InfiniBand port cluster configuration in a DB2 pureScale environment:

```
db2prereqcheck -v 10.1.0.0 -p -t SINGLE_IB_PORT_CLUSTER
```

```
Checking DB2 prerequisites for DB2 database version 10.1.0.0 on
operating system "Linux"
```

```
Validating Linux distribution ...
Required minimum operating system distribution: "RHEL";
Version: "5"; Service pack: "6".
Actual operating system distribution Version: "6";
Service pack: "1".
Requirement matched.
```

```
...
```

```
Validating Infiniband Support Package: libibcm.x86_64 ...
Package (or file) found: libibcm.x86_64
Requirement matched.
```

```
Validating Infiniband Support Package: librdmacm.x86_64 ...
Package (or file) found: librdmacm.x86_64
Requirement matched.
```

```
...
```

```
Validating Reliable Scalable Cluster Technology Package:
librdmacm.i686 ...
Package (or file) found: librdmacm.i686
Requirement matched.
```

```
Validating Reliable Scalable Cluster Technology Package:
libcxgb3.i686 ...
Package (or file) found: libcxgb3.i686
Requirement matched.
```

```
...
```

```
DBT3533I The db2prereqcheck utility has confirmed that
all installation prerequisites were met for DB2 database
server with DB2 pureScale feature. Version: "10.1.0.0".
```

On supported Linux operating systems, to validate for a multiple RoCE port cluster configuration in a DB2 pureScale environment:

```
db2prereqcheck -v 10.1.0.0 -p -t MULTI_ROCE_PORT_CLUSTER
```

```
=====
Checking DB2 prerequisites for DB2 database version 10.1.0.0 on operating
system "Linux"
```

```
Validating Linux distribution ...
```

```

Required minimum operating system distribution: "RHEL"; Version: "6";
Service pack: "1".
Actual operating system distribution Version: "6"; Service pack: "1".
Requirement matched.

...

Validating Infiniband Support Package: libibcm.x86_64 ...
Package (or file) found: libibcm.x86_64
Requirement matched.

Validating Infiniband Support Package: librdmacm.x86_64 ...
Package (or file) found: librdmacm.x86_64
Requirement matched.

...

Validating High Performance Networking Package: libibverbs-rocee.x86_64 ...
Package (or file) found: libibverbs-rocee.x86_64
Requirement matched.

Validating High Performance Networking Package: libmlx4-rocee.x86_64 ...
Package (or file) found: libmlx4-rocee.x86_64
Requirement matched.

Validating Reliable Scalable Cluster Technology Package: libibcm.i686 ...
Package (or file) found: libibcm.i686
Requirement matched.

Validating Reliable Scalable Cluster Technology Package: librdmacm.i686 ...
Package (or file) found: librdmacm.i686
Requirement matched.

...

DBT3533I The db2prereqcheck utility has confirmed that all installation
prerequisites were met for DB2 database server with DB2 pureScale
feature. Version: "10.1.0.0".

Sample db2prereqcheck command output on Linux operating system. For DB2
pureScale uDAPL only, use the -u option.
> db2prereqcheck -u -v 10.1.0.0
=====
Checking DB2 Prerequisites for DB2 Version 10.1.0.0 on Linux

Validating ofed ...
Required minimum version and release for ofed: 1.5.2-(null)
Requirement matched.

Validating sg3_utils ...
Package (or file) found: sg3_utils
Requirement matched.

Validating sg_persist ...
Package (or file) found: /usr/bin/sg_persist
Requirement matched.

All requirement matched for DB2 Server version 10.1.0.0 .

=====
Checking DB2 Prerequisites for DB2 Version 10.1.0.0 on Linux

Validating ofed ...
Required minimum version and release for ofed: 1.5.2-(null)
Requirement matched.

```

```
Validating sg3_utils ...  
Package (or file) found: sg3_utils  
Requirement matched.
```

```
Validating sg_persist ...  
Package (or file) found: /usr/bin/sg_persist  
Requirement matched.
```

All requirement matched for DB2 Server version 10.1.0.0 DB2 pureScale feature.

Sample **db2prereqcheck** command output on Linux operating system. For client installation, use the **-c** option.

```
> db2prereqcheck -c -v 10.1.0.0
```

```
=====
```

Checking DB2 Prerequisites for DB2 Version 10.1.0.0 on Linux

```
Validating Linux distribution ...  
Required minimum SUSE version : (null) SP (null)
```

```
Validating sles-release ...  
Required minimum version and release for sles-release: 10-15.45.8  
Actual version of package: 11.1
```

Requirement matched.

```
Validating kernel level ...  
Required minimum kernel level : 2.6.16  
Actual kernel level: 2.6.32.45  
Requirement matched.
```

```
Validating C++ Library version ...  
Required minimum C++ library: libstdc++.so.6  
Standard C++ library is located at: /usr/lib64/libstdc++.so.6.0.10  
Actual C++ library: CXXABI_1.3.1
```

```
Validating 32 bit version of libstdc++.so.6 ...  
Found the 64 bit /usr/lib64/libstdc++.so.6 at /usr/lib64  
Found the 32 bit /usr/lib/libstdc++.so.6 at /usr/lib  
Requirement matched.
```

All requirement matched for DB2 Client version 10.1.0.0 .

```
=====
```

Checking DB2 Prerequisites for DB2 Version 10.1.0.0 on Linux

```
Validating Linux distribution ...  
Required minimum SUSE version : (null) SP (null)
```

```
Validating sles-release ...  
Required minimum version and release for sles-release: 10-15.45.8  
Actual version of package: 11.1
```

Requirement matched.

```
Validating kernel level ...  
Required minimum kernel level : 2.6.16  
Actual kernel level: 2.6.32.45  
Requirement matched.
```

```
Validating C++ Library version ...  
Required minimum C++ library: libstdc++.so.6  
Standard C++ library is located at: /usr/lib64/libstdc++.so.6.0.10  
Actual C++ library: CXXABI_1.3.1
```

```
Validating 32 bit version of libstdc++.so.6 ...
```

```

Found the 64 bit /usr/lib64/libstdc++.so.6 at /usr/lib64
Found the 32 bit /usr/lib/libstdc++.so.6 at /usr/lib
Requirement matched.

Validating libc.so version ...
glibc library is located at: /lib64/libc-2.11.1.so
Required minimum glibc library version: 2.4.0
Actual glibc library version: 2.11.1
Requirement matched.

Validating ofed ...
Required minimum version and release for ofed: 1.5.2-(null)
Requirement matched.

Validating sg3_utils ...
Package (or file) found: sg3_utils
Requirement matched.

Validating sg_persist ...
Package (or file) found: /usr/bin/sg_persist
Requirement matched.

All requirement matched for DB2 Client version 10.1.0.0 DB2 pureScale
feature.

Sample db2prereqcheck command output on Linux operating system. To display
validation summary only, use the -s option.
> db2prereqcheck -s -v 10.1.0.0
All requirement matched for DB2 Server version 10.1.0.0 .

All requirement matched for DB2 Server version 10.1.0.0 DB2 pureScale
feature.

```

Part 2. DB2 servers and IBM data server clients

A DB2 database system consists of a DB2 server and IBM data server clients.

A DB2 server is a relational database management system (RDBMS) that delivers data to its IBM data server clients. If you plan to use a database that resides on this computer, install a DB2 server. For more information about DB2 servers, go to <http://www.ibm.com/software/data/db2/>

An IBM data server client is an application that allows you to run commands and SQL statements against a DB2 server, connect to a remote DB2 server, and access its databases.

Chapter 2. An overview of installing DB2 database servers (Linux and UNIX)

This topic outlines the steps for installing your DB2 server product on AIX, HP-UX, Linux, and Solaris.

Procedure

To install your DB2 server product:

1. Review your DB2 product prerequisites.
2. Review DB2 upgrade information if applicable.
3. Modify kernel parameters on HP-UX, Linux, and Solaris. On all platforms, except for Linux on x86_32, you must install a 64-bit kernel before proceeding with the installation, otherwise the installation will fail.
4. Prepare the installation media:

Product DVD

If the DB2 product DVD does not automount, mount your DB2 product DVD.

Installation image

If you downloaded an installation image, untar the file.

5. Install your DB2 product using one of the available methods:
 - The DB2 Setup wizard
 - A silent installation using a response file
 - Payload file deployment

For DB2 servers, you can use the DB2 Setup wizard to perform installation and configuration tasks, such as:

- Selecting DB2 installation type (typical, compact, or custom).
 - Selecting DB2 product installation location.
 - Install the languages that you can specify later as the default language for the product interface and messages.
 - Install or upgrade the IBM Tivoli® System Automation for Multiplatforms (Linux and AIX).
 - Setting up a DB2 instance.
 - Setting up the DB2 Administration Server (including DAS user setup).
 - Setting up the DB2 Text Search server.
 - Setting up Administration contact and health monitor notification.
 - Setting up and configuring your instance setup and configuration (including instance user setup).
 - Setting up Informix® data source support.
 - Preparing the DB2 tools catalog.
 - Specify the DB2 Information Center port.
 - Creating response files.
6. If you installed a DB2 server using a method other than the DB2 Setup wizard, post-installation configuration steps are required.

Chapter 3. An overview of installing DB2 database servers (Windows)

This topic outlines the steps for installing your DB2 server product on Windows.

Procedure

To install your DB2 server product:

1. Review your DB2 database product prerequisites.
2. Review DB2 database upgrade information if applicable.
3. Prepare the installation media:

Product DVD

Insert the DB2 database product DVD into your DVD-ROM drive.

Installation image

If you downloaded an installation image, extract the file.

4. Install your DB2 database product using:

- the DB2 Setup wizard.
- a silent installation with a response file.

You can use the DB2 Setup wizard to perform installation and configuration tasks, such as:

- Selecting DB2 database installation type (typical, compact, or custom).
- Selecting DB2 database product installation location.
- Setting up IBM SSH Server for Windows
- Installing the languages that you can specify later as the default language for the product interface and messages.
- Setting up DB2 Extended Security.
- Setting up a DB2 instance (including instance user setup. By default, it is same as the DAS user).
- Setting up the DB2 Administration Server (including DAS user setup).
- Setting up the DB2 Text Search server.
- Setting up Administration contact and health monitor notification.
- Setting up and configuring your instance setup and configuration (including instance user setup).
- Preparing the DB2 tools catalog.
- Specifying the *DB2 Information Center* port.
- Creating response files.
- Installing additional products.

Chapter 4. Non-root installation overview (Linux and UNIX)

The DB2 installer automatically creates and configures a non-root instance during a non-root installation. As a non-root user, you can customize the configuration of the non-root instance during the installation. You can also use and maintain the installed DB2 database product without root privileges.

The non-root installation of a DB2 database product has one DB2 instance with most features enabled by default.

A non-root installation can be attractive for many groups, such as:

- Enterprises that have thousands of workstations and users who want to install a DB2 database product without consuming a system administrator's time
- Application developers who are not typically system administrators but use DB2 database products to develop applications
- Independent Software Vendors (ISVs) who develop software that does not require root user authority yet embeds a DB2 database product

Although non-root installations have most of the function of root installations, there are some differences and limitations. You can lift some of the limitations by having a root user run the **db2rfe** command.

Installing DB2 database servers as a non-root user

Most DB2 database products can be installed as a non-root user.

Before you begin

Before you install any DB2 database product as a non-root user, be aware of the differences between root installations and non-root installations, and the limitations of non-root installations. For more information on non-root installation, see “Non-root installation overview (Linux and UNIX)”.

Prerequisites for installing a DB2 database product as a non-root user are:

- You must be able to mount the installation DVD, or have it mounted for you.
- You must have a valid user ID that can be used as the owner of a DB2 instance.

User IDs have the following restrictions and requirements:

- Must have a primary group other than guests, admins, users, and local
- Can include lowercase letters (a-z), numbers (0-9), and the underscore character (_)
- Cannot be longer than eight characters
- Cannot begin with IBM, SYS, SQL, or a number
- Cannot be a DB2 reserved word (USERS, ADMINS, GUESTS, PUBLIC, or LOCAL), or an SQL reserved word
- Cannot use any User IDs with root privilege for the DB2 instance ID, DAS ID or fenced ID.
- Cannot include accented characters
- If existing user IDs are specified instead of creating new user IDs, make sure that the user IDs:

- Are not locked
 - Do not have expired passwords
 - The hardware and software prerequisites that exist for the product you are installing apply to the non-root user just as they do for root users.
 - On AIX Version 5.3, Asynchronous I/O (AIO) must be enabled. It is strongly recommended the system has I/O Completion Ports (IOCP) enabled.
 - Your home directory must be a valid DB2 path.
- DB2 installation paths have the following rules:
- Can include lowercase letters (a-z), uppercase letters (A-Z), and the underscore character (_)
 - Cannot exceed 128 characters
 - Cannot contain spaces
 - Cannot contain non-English characters

About this task

Installing DB2 database products as a non-root user is transparent to the non-root user. In other words, there is nothing special a non-root user needs to do to install a DB2 database product, other than being logged in as a non-root user.

Procedure

To perform a non-root installation:

1. Log in as a non-root user
2. Install your DB2 database product using any of the methods available to you.
Options include:
 - The DB2 Setup wizard (GUI install)
 - The **db2setup** command with a response file (silent install)

Note: Since non-root users cannot choose the directory where DB2 database products are installed, any **FILE** keyword in your response file is ignored.

3. After the DB2 database product is installed, you must open a new login session to use the non-root DB2 instance. Alternatively, you can use the same login session if you set up the DB2 instance environment with `$HOME/sql1lib/db2profile` (for Bourne shell and Korn shell users) or `$HOME/sql1lib/db2chsrc` (for C shell users), where `$HOME` is the non-root user's home directory.

What to do next

After the DB2 database product is installed, verify your operating system user process resource limits (ulimits). If the minimum ulimit values are not met, the DB2 engine can encounter unexpected operating resource shortage errors. These errors can lead to a DB2 database system outage.

Enabling root-based features in non-root installations with db2rfe

There are several features and abilities in non-root installations that are initially unavailable but can be enabled by running the **db2rfe** command.

Before you begin

This task requires root user authority.

Procedure

To enable the features and abilities that are initially unavailable in non-root installations:

1. Locate the sample configuration files. Two sample configuration files are provided:
 - `$HOME/sql11ib/instance/db2rfe.cfg` is pre-configured with default values for the non-root DB2 instance
 - `$HOME/sql11ib/cfg/db2rfe.cfg.sample` is not configuredwhere `$HOME` is the non-root user's home directory.
2. Copy one of the sample configuration files to a different location so the original file remains unaltered.
3. Update the copied configuration file as needed. This configuration file is input to the **db2rfe** command.

An example of a configuration file is:

```
INSTANCENAME=db2inst2
SET_ULIMIT=NO
ENABLE_HA=NO
ENABLE_OS_AUTHENTICATION=NO
RESERVE_REMOTE_CONNECTION=NO
  **SVCENAME=db2c_db2inst2
  **SVCEPORT=48000
RESERVE_TEXT_SEARCH_CONNECTION=NO
  **SVCENAME_TEXT_SEARCH=db2j_db2inst2
  **SVCEPORT_TEXT_SEARCH=55000
```

Note:

- The value for the **INSTANCENAME** parameter is filled in automatically by DB2 installer
- The **SET_ULIMIT** parameter is available only on AIX. On other operating systems, a user with root authority needs to set ulimit values manually.
- The default value for the other keywords is NO
- Child parameters (such as **SVCENAME**) are commented out by default. Comments are denoted with ******
- If you set a parameter to YES, and if it has any child parameters, it is recommended that you uncomment the child parameters and provide appropriate values. Any port values that are provided are examples. Ensure the port values you assign are free.

The following example is provided to show an edited configuration file that will enable the features and abilities mentioned in the following list:

- High availability
- Operating system-based authentication
- DB2 Text Search, with a service name of `db2j_db2inst2` and a port value of 55000

To enable these features and abilities, edit the configuration file as follows:

```
INSTANCENAME=db2inst2
SET_ULIMIT=NO
ENABLE_HA=YES
ENABLE_OS_AUTHENTICATION=YES
```



```

RESERVE_REMOTE_CONNECTION=NO
**SVCENAME=db2c_db2inst2
**SVCEPORT=48000
RESERVE_TEXT_SEARCH_CONNECTION=YES
SVCENAME_TEXT_SEARCH=db2j_db2inst2
SVCEPORT_TEXT_SEARCH=55000

```

4. Log in with root user authority.
5. Navigate to the `$HOME/sqllib/instance` directory, where `$HOME` represents the non-root user's home directory.
6. Run the **db2rfe** command using the following syntax:
`db2rfe -f config_file`

where *config_file* is the configuration file created in Step 3 on page 19.

What to do next

To keep root-based features enabled on non-root installations, rerun the **db2rfe** command after applying fix packs or upgrading to a new version.

Differences between root installations and non-root installations

In addition to a few limitations, the directory structure of a non-root installation is slightly different than the directory structure of a root installation.

During a root installation, subdirectories and files for the DB2 database product are created in a directory of the root user's choosing.

Unlike root users, non-root users cannot choose where DB2 database products are installed. Non-root installations are always placed in the `$HOME/sqllib` directory, where `$HOME` represents the non-root user's home directory. The layout of the subdirectories within the `sqllib` directory of a non-root is similar to that of a root installation.

For root installations, multiple instances can be created. Instance ownership is associated with the user ID under which the instance was created.

Non-root installations can have only one DB2 instance. The non-root installation directory contains all of the DB2 database product files and instance files with no soft links.

The following table summarizes the differences between root installations and non-root installations.

Table 1. Differences between root installations and non-root installations

Criteria	Root installations	Non-root installations
User can select installation directory	Yes	No. DB2 database products are installed under the user's home directory.
Number of DB2 instances allowed	Multiple	One
Files deployed during installation	Program files only. Instances must be created after installation.	Program files and instance files. The DB2 database product is ready for use immediately after installation.

Table 1. Differences between root installations and non-root installations (continued)

Criteria	Root installations	Non-root installations
Upgrade version and instance	No	No need to uninstall the old version before installing the new version. Install new version, and upgrade the instance, together.

Limitations of non-root installations

In addition to the differences between root installations and non-root installations, there are several limitations on non-root installations. This topic discusses the limitations to help you decide if you want to use a non-root installation.

Product limitations

Some DB2 database products are not supported in non-root installations:

- DB2 Net Search Extender
- Locally installed *DB2 Information Center*

Note: The locally installed *DB2 Information Center* is not supported in non-root installations because it requires root user authority to start the daemon. However, a non-root-installation DB2 instance can be configured to use a locally installed *DB2 Information Center* if it is installed on the same computer.

Features and tools limitations

The following features and tools are not available in non-root installations:

- The DB2 Administration Server (DAS) and its associated commands: **dasCRT**, **dasdrop**, **daslist**, **dasmigr**, and **dasupdt**
- The ability for the **db2governor** to increase priority is not supported
- In the Work Load Manager (WLM), attempts to set agent priority in a DB2 service class in a non-root DB2 instance are allowed. However, the agent priority will not be respected, and no SQLCODE error is returned.
- Automatic starting of non-root DB2 instances at system reboot is not supported

Health monitor limitations

The following health monitor features are not supported in non-root installations:

- Running script or task actions on alert occurrences
- Sending alert notifications

Partitioned database limitation

Only single-partition databases are supported in non-root installations. You cannot add additional database partitions.

Listing DB2 database products

The output produced by the **db2ls** command, when run as a non-root user, is different than the output produced when run as a root user. For details, refer to the **db2ls** command topic.

DB2 copies

Each non-root user can have only one copy of a DB2 database product installed.

DB2 instance limitation

In non-root installations, one DB2 instance is created during installation. Additional instances cannot be created.

DB2 instance actions can be performed only by the instance owner

Root installations and non-root installations can coexist on the same computer in different installation paths. However, a non-root instance can be updated, or dropped (using the **db2_deinstall** command), only by the non-root user who owns the non-root instance.

A DB2 instance created by a user with root user authority can be updated or dropped only by a user with root user authority.

DB2 instance commands

The following DB2 instance commands are unavailable in non-root installations:

db2icrt

When installing a DB2 database product as a non-root user, a single instance is automatically created and configured. Further instances cannot be created in non-root installations. However, if the automatically created instance needs to be configured, you can use the non-root install configuration command, **db2nrcfg**.

db2iupdt

The **db2iupdt** command cannot be used for non-root instances. Instead, use the non-root install configuration command (**db2nrcfg**) to update the non-root DB2 instance. However, updating the non-root instance is normally not required because it gets updated automatically when updating your DB2 database product.

db2idrop

The instance that gets automatically created during non-root installations cannot be dropped. The DB2 database product must be uninstalled to drop the DB2 instance.

db2iupgrade

Upgrading is not supported for non-root installations. To upgrade a non-root instance, use the **db2nrcfg** command.

Upgrading limitation

Root instances cannot be upgraded to a non-root instance.

Post-installation actions can be performed only by the DB2 instance owner

Root installations and non-root installations can coexist on the same computer. However, only the original non-root user who installed the DB2 database product can perform subsequent actions such as:

- Applying fix packs
- Adding features
- Installing add-on products

Adjusting ulimit values

The **ulimit** command on UNIX and Linux operating systems sets or reports user resource limits, such as data and stack limits. For root instances, the database server dynamically updates required ulimit settings without changing the permanent settings. However, for non-root instances, the ulimit settings can only be checked during installation. A warning message is displayed if the settings are inadequate. Root user authority is required to change the ulimit settings.

Limitations that can be overcome by running db2rfe

There are further limitations on non-root installations which can be overcome by running the **db2rfe** command. The following features and abilities are initially unavailable in non-root installations:

- Operating system-based authentication
- High Availability (HA) feature
- The ability to reserve service names in the `/etc/services` file
- The ability to increase user data limits (ulimits). This ability applies only to AIX. On other operating systems, user data limits must be increased manually.

Run the Enable root features for non-root install command (**db2rfe**) to enable these features and abilities. Running the **db2rfe** command is optional, and must be run by a user with root user authority.

Authentication type in non-root installations

Operating system-based authentication is the default authentication type for DB2 database products. Since non-root installations do not support operating system-based authentication, if you choose not to run the **db2rfe** command after installing your DB2 database product as a non-root user, then you must manually set the authentication type. You can do so by updating the following parameters in the database manager configuration (`dbm cfg`) file:

- **clnt_pw_plugin** (Client userid-password plug-in configuration parameter)
- **group_plugin** (Group plug-in configuration parameter)
- **srvcon_pw_plugin** (Userid-password plug-in for incoming connections at the server configuration parameter)

Reducing the size of your DB2 product installation image

You can use the **db2iprune** command to reduce the size of a DB2 database product installation image.

About this task

This tool is useful for large-scale deployments of the DB2 database product, and for embedding DB2 within an application. The **db2iprune** utility removes the files associated with those features and languages based on an input file. The input file (`.prn` file) allows you to specify which features and languages you want to remove from the installation image. The result is a new, smaller DB2 installation image that can be installed using the regular DB2 installation methods. Reducing the size of an installation image is also referred to as pruning the installation image.

Be aware of the following restrictions before proceeding with pruning an image:

- Some components have dependency. A component can be pruned only if there is no other component not being pruned that depends on it. For example, `INFORMIX_DATA_SOURCE_SUPPORT` depends on `DB2_DATA_SOURCE_SUPPORT`. Either `INFORMIX_DATA_SOURCE_SUPPORT` or both `INFORMIX_DATA_SOURCE_SUPPORT` and `DB2_DATA_SOURCE_SUPPORT` can be pruned, but `DB2_DATA_SOURCE_SUPPORT` cannot be pruned and leave `INFORMIX_DATA_SOURCE_SUPPORT` in the image.
- Products can be removed, but at least one product must be left.

- English language cannot be removed.
- On Linux and UNIX operating systems, a language can be pruned in either of the following ways:
 - you can run the **db2iprune** command from the DB2 National Language Pack (NLPACK), where the NLPACK can be anywhere. In this case, only languages can be pruned from inside the NLPACK.
 - you can run the **db2iprune** command from the DB2 database product DVD. In this case, if NLPACK is a subdirectory, you can prune languages, and, products or components or both, in one execution of the **db2iprune** command.

Procedure

To reduce the size of your DB2 database product installation image:

1. Create a customized input file using the sample input file (.prn file) located in the `db2/plat/utilities/db2iprune/` directory on the product DVD or downloaded image as a base. You can remove the * to uncomment that particular line. You can also specify these keywords:

PRUNE_PROD

Specifies the DB2 database product to remove. This keyword is optional. More than one DB2 database product can be removed at a time, but at least one product must remain in the installation image. For example, use

```
PRUNE_PROD          = CLIENT
```

to prune the IBM Data Server Client.

PRUNE_COMP

Specifies the DB2 component to remove. This keyword is optional. More than one DB2 component can be removed at a time. When a component is removed, the specified component is removed from all applicable products. For example, use

```
PRUNE_COMP          = FIRST_STEPS
```

to prune First Steps.

PRUNE_LANG

Specifies the installed language to remove. The English language is mandatory and cannot be removed. This keyword is optional. More than one language can be removed at a time. All languages except English are removed by specifying **PRUNE_LANG=all**. When a language is removed, the specified language is removed from all applicable products. For example, use

```
PRUNE_LANG          = CZ
```

to prune the Czech language.

PRUNE_TSAMP

Linux and UNIX operating systems only. Specifies that IBM Tivoli System Automation for Multiplatforms (SA MP) is to be removed from the installation image.

PRUNE_VSAI

On Windows operating systems only. Specifies that IBM Database Add-Ins for Visual Studio is to be removed from the installation image.

2. From the command line, run the **db2iprune** command. For information about the command parameters, see “db2iprune - Reduce installation image size command”.

Results

You can use any of the regular DB2 installation methods to install and maintain a pruned DB2 installation image:

DB2 Setup Wizard installation

For a typical installation, the regular typical components for that product are installed without the components removed by the **db2iprune** command.

For a compact installation, the regular compact components for that product are installed without the components removed by the **db2iprune** command.

For a custom installation, only the remaining components are displayed in the feature selection panel. The components removed by the **db2iprune** command are not displayed as optional components to install. However, on Linux and UNIX operating systems, the removed languages will still be displayed in the language selection panel. In this case, ensure that you do not select a language that has been removed from the image using the **db2iprune** command; if you select a language that has been removed, you will receive an error message.

Response file installation

If you plan to use a response file for an unattended installation, ensure that you specify only the languages and features available in the DB2 pruned installation image. If you select a component that has been removed, you will get an error message.

Fix pack installation

Since Windows DB2 fix packs are full installation images, the **db2iprune** command can be used with fix pack images. The fix pack application process is the same for full and pruned images. When the DB2 fix pack is installed, it detects and updates only the components that were installed and ignores any components that are not installed. If the **db2iprune** command is used with a fix pack image, ensure that the fix pack image contains all of the components that were initially installed. If the fix pack image does not contain all the installed components, you will receive an error about missing files when the fix pack application is attempted.

Chapter 5. Java software support for DB2 database products

You require the appropriate level of IBM Software Development Kit (SDK) for Java to use Java-based tools and to create and run Java applications, including stored procedures and user-defined functions.

If the IBM SDK for Java is required by a component being installed and the SDK for Java is not already installed in that path, the SDK for Java will be installed if you use either the DB2 Setup wizard or a response file to install the product.

The SDK for Java is not installed with IBM Data Server Runtime Client or IBM Data Server Driver Package.

The following table lists the installed SDK for Java levels for DB2 database products according to operating system platform:

Operating System Platform	SDK for Java level
AIX	SDK 7
HP-UX for Itanium-based systems	SDK 6
Linux on x86	SDK 7
Linux on AMD64/EM64T	SDK 7
Linux on zSeries®	SDK 7
Linux on POWER®	SDK7
Solaris Operating System	SDK 7
Windows x86	SDK 7
Windows x64	SDK 7

Note:

1. The SDK for Java software can be downloaded from the developerWorks® Web page at: <http://www.ibm.com/developerworks/java/jdk/index.html> . For a list of the supported levels of the SDK for Java, see the table later in this section entitled DB2 for Linux, UNIX, and Windows support for SDKs for Java.

Note: For Windows operating system platforms, use the IBM Development Package for Eclipse downloads.

2. DB2 GUI tools only run on Linux on x86, Linux on AMD64/EM64T, Windows x86, and Windows x64.
3. On Windows x86 and Linux on x86:
 - the 32-bit SDK is installed
 - 32-bit applications and Java external routines are supported
4. On all supported platforms (except Windows x86, and Linux on x86):
 - 32-bit applications are supported
 - 32-bit Java external routines are not supported
 - 64-bit applications and Java external routines are supported

Supported Java application development software

The following table lists the supported levels of the SDK for Java. The listed levels and forward-compatible later versions of the same levels are supported.

Because there are frequent SDK for Java fixes and updates, not all levels and versions have been tested. If your database application has problems that are related to the SDK for Java, try the next available version of your SDK for Java at the given level.

Versions of SDK for Java, other than IBM SDK, are supported only for building and running stand-alone Java applications. For building and running new Java stored procedures and user-defined functions, only the IBM SDK for Java that is included with the DB2 for Linux, UNIX, and Windows product is supported. For running Java stored procedures and user-defined functions that were built by prior DB2 releases, refer to Table 1, column "Java Stored Procedures and User Defined Functions" for details.

Table 2. DB2 for Linux, UNIX, and Windows supported levels of SDKs for Java

	Java applications using JDBC driver db2java.zip or db2jcc.jar	Java applications using JDBC driver db2jcc4.jar ⁷	Java Stored Procedures and User Defined Functions	DB2 Graphical Tools
AIX	1.4.2 to 7	6 and 7	1.4.2 ⁶ to 7 ⁵	N/A
HP-UX for Itanium-based systems	1.4.2 to 6 ¹	6 ¹	1.4.2 ⁶ to 6	N/A
Linux on POWER	1.4.2 to 7 ^{3,4}	6 and 7 ^{3,4}	1.4.2 ⁶ to 7	N/A
Linux on x86	1.4.2 to 7 ^{2,3,4}	6 and 7 ^{2,3,4}	1.4.2 ⁶ to 7	5 to 7
Linux on AMD64 and Intel EM64T processors	1.4.2 to 7 ^{2,3,4}	6 and 7 ^{2,3,4}	1.4.2 ⁶ to 7	N/A
Linux on zSeries	1.4.2 to 7 ^{3,4}	6 and 7 ^{3,4}	1.4.2 ⁶ to 7	N/A
Solaris operating system	1.4.2 to 7 ²	6 and 7 ²	1.4.2 ⁶ to 7	N/A
Windows on x86	1.4.2 to 7 ²	6 and 7 ²	1.4.2 ⁶ to 7	5 to 7
Windows on x64, for AMD64 and Intel EM64T processors	1.4.2 to 7 ²	6 and 7 ²	1.4.2 ⁶ to 7	5 to 7

Note:

1. The same levels of the SDK for Java that are available from Hewlett-Packard are supported for building and running stand-alone client applications that run under the IBM Data Server Driver for JDBC and SQLJ.
2. The same levels of the SDK for Java that are available from Oracle are supported for building and running stand-alone applications with the IBM Data Server Driver for JDBC and SQLJ. However, if you set the IBM Data Server Driver for JDBC and SQLJ property securityMechanism for a type of security that uses encryption, the SDK for Java must support the type of encryption that you use. For example, the SDK for Java that you use might support 256-bit AES (strong) encryption, but not 56-bit DES (weak) encryption. You can specify the encryption algorithm by setting the IBM Data Server Driver

for JDBC and SQLJ property encryptionAlgorithm. To use 256-bit AES encryption, set encryptionAlgorithm to 2. When you use 256-bit AES encryption with the SDK for Java from Oracle, you might need to install the JCE Unlimited Strength Jurisdiction Policy File, which is available from Oracle.

3. A minimum level of SDK for Java 1.4.2 SR6 is required for SUSE Linux Enterprise Server (SLES) 10. A minimum level of SDK for Java 1.4.2 SR7 is required for Red Hat Enterprise Linux (RHEL) 5.
4. SDK for Java 6 support on Linux requires SDK for Java 6 SR3 or later.
5. If SDK for Java 6 SR2 or later is used, set DB2LIBPATH=*java_home*/jre/lib/ppc64.
6. Support for Java stored procedures and user-defined functions built by IBM SDK for Java 1.4.2 was deprecated in Version 9.7 and might be removed in a future release. IBM SDK for Java 1.4.2 has an End of Service date of September 2011. It is recommended to remove SDK for Java 1.4.2 dependency well before this date. Removing this dependency can be done by rebuilding Java stored procedures and user-defined functions with the SDK for Java included in DB2 Version 9.1, DB2 Version 9.5, DB2 Version 9.7 or DB2 Version 10.1 .
7. Java 6 is sufficient if you need to use JDBC 4.0 functions only. Java 7 is required if you need to use JDBC 4.1 functions.

Chapter 6. Preparing to install DB2 database servers

Before installing DB2 database server, ensure that the necessary prerequisites are met, such as disk, memory, and paging space requirements. There are also additional prerequisites that depend on your operating system.

You can also install multiple DB2 copies on the same computer. For Windows systems, there is a difference between installing one or multiple DB2 copies. Each DB2 copy can be at the same or different code levels. A DB2 copy is a group of DB2 products that are installed at the same location. For Linux and UNIX systems, each DB2 copy can be at the same or different code levels. Root installation of DB2 products can be installed to an installation path of your choice.

Disk and memory requirements

Ensure that an appropriate amount of disk space is available for your DB2 environment, and allocate memory accordingly.

Disk requirements

The disk space required for your product depends on the type of installation you choose and the type of file system you have. The DB2 Setup wizard provides dynamic size estimates based on the components selected during a typical, compact, or custom installation.

Remember to include disk space for required databases, software, and communication products. Ensure that the file system is not mounted with concurrent I/O (CIO) option.

On Linux and UNIX operating systems, 2 GB of free space in the /tmp directory is recommended, and at least 512 MB of free space in the /var directory is required.

Note: On Linux and UNIX operating systems, you must install your DB2 product in an empty directory. If the directory that you have specified as the install path contains subdirectories or files, your DB2 installation might fail.

On Windows operating systems the following free space is recommended in addition to that of your DB2 product:

- 40 MB in the system drive
- 60 MB in the temporary folder specified by the temp environment variable.

Memory requirements

Memory requirements are affected by the size and complexity of your database system, the extent of database activity, and the number of clients accessing your system. At a minimum, a DB2 database system requires 256 MB of RAM¹. For a system running just a DB2 product and the DB2 GUI tools, a minimum of 512 MB of RAM is required. However, 1 GB of RAM is recommended for improved performance. These requirements do not include any additional memory requirements for other software that is running on your system. For IBM data

1. DB2 products that run on HP-UX Version 11i for Itanium-based systems require a minimum of 512 MB of RAM.

server client support, these memory requirements are for a base of five concurrent client connections. For every additional five client connections, an additional 16 MB of RAM is required.

For DB2 server products, the self-tuning memory manager (STMM) simplifies the task of memory configuration by automatically setting values for several memory configuration parameters. When enabled, the memory tuner dynamically distributes available memory resources among several memory consumers including sort, the package cache, the lock list, and buffer pools.

Paging space requirements

DB2 requires paging, also called swap to be enabled. This configuration is required to support various functions in DB2 which monitor or depend on knowledge of swap/paging space utilization. The actual amount of swap/paging space required varies across systems and is not solely based on memory utilization by application software. It is only strictly required by DB2 on the Solaris and HP platforms due to their use of early paging space allocation.

A reasonable minimum swap/paging space configuration for most systems is 25-50% of RAM. Solaris and HP systems with many small databases or multiple databases tuned by STMM might require a paging space configuration of 1 x RAM or higher. These higher requirements are due to virtual memory pre-allocated per database / instance, and retained virtual memory in the case of STMM tuning multiple databases. Additional swap/paging space might be wanted to provision for unanticipated memory overcommitment on a system.

Installation prerequisites for database servers

Before you install your DB2 database product, ensure that the system you choose meets the necessary operating system, hardware, software, communications, disk, and memory requirements.

There are different prerequisites for AIX, UNIX, HP-UX, Linux, Solaris, and Windows operating systems.

Prerequisites for a DB2 database server installation (Linux and UNIX)

Before you install your DB2 database product, ensure that the prerequisites for each of the tasks are met for installing on Linux or UNIX operating systems.

DB2 users and groups (Linux and UNIX)

The DB2 Setup wizard creates the users and groups automatically (if needed) during the installation of your DB2 database product.

Note: This topic does not apply to non-root installations.

If you are using the DB2 Setup wizard, you can create the following users and groups during installation. To manually create the following users and groups, see “Creating group and user IDs for a DB2 database installation (Linux and UNIX)” on page 35. Three users and three groups are used to operate DB2 on Linux and UNIX operating systems.

Instance owner

The DB2 instance is created in the instance owner home directory. This user ID controls all DB2 processes and owns all filesystems and devices

used by the databases contained within the instance. The default user is db2inst1 and the default group is db2iadm1.

When using the DB2 Setup wizard, the default action is to create a new user for your DB2 instance. The default name is db2inst1. If that user name already exists, the DB2 Setup wizard searches through user names (db2inst2, db2inst3, and so on). The search continues until a user name is identified that is not already an existing user on the system as the default instance owner ID. If you choose to proceed this user is created by the DB2 Setup wizard. However, you also have a choice to specify any existing user as the instance owner.

This method for user name creation also applies to the creation of fenced users and DB2 administration server users.

Fenced user

The fenced user is used to run user defined functions (UDFs) and stored procedures outside of the address space used by the DB2 database. The default user is db2fenc1 and the default group is db2fadm1. If you do not need this level of security, for example in a test environment, you can use your instance owner as your fenced user.

DB2 administration server user

The user ID for the DB2 administration server user is used to run the DB2 administration server (DAS) on your system. The default user is dasusr1 and the default group is dasadm1.

There is only one DAS per computer. One DAS services one or more database instances, including database instances that belong to different installations. The DAS can service database instances whose release level is lower than the release level of the DAS. However, for database instances whose release level is higher than the release level of the DAS, the DAS must be migrated to a higher level. The DAS release level must be as high (or higher) than the release level of any of the database instances it services.

Important: The DB2 Administration Server (DAS) has been deprecated in Version 9.7 and might be removed in a future release. The DAS is not supported in DB2 pureScale environments. Use software programs that use the Secure Shell protocol for remote administration. For more information, see “DB2 administration server (DAS) has been deprecated” at .

User ID restrictions

User IDs have the following restrictions and requirements:

- Must have a primary group other than guests, admins, users, and local
- Can include lowercase letters (a-z), numbers (0-9), and the underscore character (_)
- Cannot be longer than eight characters
- Cannot begin with IBM, SYS, SQL, or a number
- Cannot be a DB2 reserved word (USERS, ADMINS, GUESTS, PUBLIC, or LOCAL), or an SQL reserved word
- Cannot use any User IDs with root privilege for the DB2 instance ID, DAS ID or fenced ID.
- Cannot include accented characters

- If existing user IDs are specified instead of creating new user IDs, make sure that the user IDs:
 - Are not locked
 - Do not have expired passwords

Fast communications manager (Linux and UNIX)

The fast communications manager (FCM) provides communications support for partitioned database environments.

In multiple member environments, each member has a pair of FCM daemons to support communication between members that is related to agent requests. One daemon is for sending communications, and the other is for receiving. These daemons and supporting infrastructure are activated when an instance is started. FCM communication is also used for agents working within the same member; this type of communication is also known as intra-member communication.

The FCM daemon collects information about communication activities. You can obtain information about FCM communications by using the database system monitor. If communications fail between members or if they re-establish communications, the FCM daemons update monitor elements with this information. The FCM daemons also trigger the appropriate action for this event. An example of an appropriate action is the rollback of an affected transaction. You can use the database system monitor to help you set the FCM configuration parameters.

You can specify the number of FCM message buffers by using the **fcm_num_buffers** database manager configuration parameter. You can specify the number of FCM channels by using the **fcm_num_channels** database manager configuration parameter. By default, the **fcm_num_buffers** and **fcm_num_channels** database manager configuration parameters are set to AUTOMATIC. If the setting is AUTOMATIC, which is the recommended setting, the FCM monitors resource usage and adjusts resources to accommodate workload demand.

Centralized user-management considerations (Linux and UNIX)

In environments that include security software, there are some installation considerations.

Note: The DB2 installation cannot update or create users and groups if they are controlled outside of the operating system. For example, LDAP can be used to control users and groups outside of the operating system.

Note: Network Information Services (NIS) and Network Information Services Plus (NIS+) features are deprecated starting with DB2 Version 9.1 Fix Pack 2. Support for these features might be removed in a future release. Lightweight Directory Access Protocol (LDAP) is the recommended solution for centralized user-management services.

At instance creation, without a security component present, the instance owner's group list is modified to include that of the database administrative server (DAS) user's primary group, if the DAS is created. If the instance creation program is unable to modify these properties, it reports that it could not. The warning message provides the necessary information to manually make the changes.

These considerations hold true for any environment in which an external security program does not allow the DB2 installation or instance creation programs to modify user characteristics.

Preparing to install DB2 for Linux on zSeries

To install a DB2 database product on an IBM zSeries that is running Linux, you must make the installation image accessible to the Linux operating system.

Before you begin

You have already obtained your DB2 database product installation image.

Procedure

- Using FTP to access the installation image
From the IBM zSeries computer running Linux:
 1. Enter the following command: `ftp yourserver.com`
where *yourserver.com* represents the FTP server where the DB2 database product installation image resides.
 2. Enter your user ID and password.
 3. Enter the following commands:

```
bin
get product_file
```

where *product_file* represents the appropriate product package name.
- Using the DB2 database product DVD over NFS to access the installation image
 1. Mount the appropriate product DVD.
 2. Export the directory where you mounted the DVD. For example, if you mounted the DVD under `/db2dvd`, then export the `/db2dvd` directory.
 3. On the IBM zSeries computer running Linux, NFS mount this directory using the following command:

```
mount -t nfs -o ro nfsservername:/db2dvd /local_directory_name
```

where *nfsservername* represents the host name of the NFS server, *db2dvd* represents the name of the directory being exported on the NFS server, and *local_directory_name* represents the name of the local directory.
 4. From the IBM zSeries computer running Linux, change to the directory where the DVD is mounted. You can do this by entering the `cd /local_directory_name` command, where *local_directory_name* represents the mount point of your product DVD.

Creating group and user IDs for a DB2 database installation (Linux and UNIX)

The DB2 Setup wizard will create these users and groups for you during the installation process. If you want, you can create them ahead of time.

Before you begin

To perform this task, you must have root user authority to create users and groups.

About this task

Three users and groups are required.

The user and group names used in the following instructions are documented in the following table. You can specify your own user and group names if they adhere to system naming rules and DB2 naming rules.

The user IDs you create will be required to complete subsequent setup tasks.

Table 3. Default users and groups

User	Example user name	Example group name
Instance owner	db2inst1	db2iadm1
Fenced user	db2fenc1	db2fadm1
DB2 administration server user	dasusr1	dasadm1

- The instance owner home directory is where the DB2 instance will be created.
- The fenced user is used to run user defined functions (UDFs) and stored procedures outside of the address space used by the DB2 database.
- The user ID for the *DB2 administration server user* is used to run the DB2 administration server on your system.

Procedure

To create the required groups and user IDs for DB2 database systems:

1. Log in as a user with root user authority.
2. Enter the appropriate commands for your operating system.

Note: These command line examples do not contain passwords. They are examples only. You can use the `passwd username` command from the command line to set the password.

AIX operating systems

To create groups on AIX, enter the following commands:

```
mkgroup id=999 db2iadm1
mkgroup id=998 db2fadm1
mkgroup id=997 dasadm1
```

Create users for each group:

```
mkuser id=1004 pgrp=db2iadm1 groups=db2iadm1
home=/home/db2inst1 db2inst1
mkuser id=1003 pgrp=db2fadm1 groups=db2fadm1
home=/home/db2fenc1 db2fenc1
mkuser id=1002 pgrp=dasadm1 groups=dasadm1
home=/home/dasusr1 dasusr1
```

Set initial password:

```
passwd db2inst1
passwd db2fenc1
passwd dasusr1
```

HP-UX operating systems

To create groups on HP-UX, enter the following commands:

```
groupadd -g 999 db2iadm1
groupadd -g 998 db2fadm1
groupadd -g 997 dasadm1
```

Create users for each group:

```
useradd -g db2iadm1 -d /home/db2inst1 -m db2inst1
useradd -g db2fadm1 -d /home/db2fenc1 -m db2fenc1
useradd -g dasadm1 -d /home/dasusr1 -m dasusr1
```

Set initial password:

```
passwd db2inst1
passwd db2fenc1
passwd dasusr1
```

Linux operating systems

To create groups on Linux operating systems, enter the following commands:

```
groupadd -g 999 db2iadm1
groupadd -g 998 db2fadm1
groupadd -g 997 dasadm1
```

Create users for each group:

```
useradd -u 1004 -g db2iadm1 -m -d /home/db2inst1 db2inst1
useradd -u 1003 -g db2fadm1 -m -d /home/db2fenc1 db2fenc1
useradd -u 1002 -g dasadm1 -m -d /home/dasusr1 dasusr1
```

Set initial password:

```
passwd db2inst1
passwd db2fenc1
passwd dasusr1
```

Solaris operating systems

To create groups on Solaris, enter the following commands:

```
groupadd -g 999 db2iadm1
groupadd -g 998 db2fadm1
groupadd -g 997 dasadm1
```

Create users for each group:

```
useradd -g db2iadm1 -u 1004 -d /export/home/db2inst1 -m db2inst1
useradd -g db2fadm1 -u 1003 -d /export/home/db2fenc1 -m db2fenc1
useradd -g dasadm1 -u 1002 -d /export/home/dasusr1 -m dasusr1
```

Set initial password:

```
passwd db2inst1
passwd db2fenc1
passwd dasusr1
```

Creating an instance using db2icrt

A DB2 instance is an environment in which you store data and run applications. Use the **db2icrt** command to create an instance.

Before you begin

On Linux or UNIX operating systems, you must have root user authority.

Note: If the DB2 fault monitor is turned on, the DB2 instance is started automatically when the **db2icrt** command finishes running. You can stop the instance by using the **db2stop** command.

Procedure

To create an instance using **db2icrt**:

1. Log in with proper authority.
2. Run the **db2icrt** command. For example, on Linux or UNIX operating systems:

```
DB2DIR/instance/db2icrt -a AuthType -u FencedID InstName
```

where:

DB2DIR

is the DB2 installation directory.

- On AIX, HP-UX, or Solaris operating systems, the default DB2 installation directory is `/opt/IBM/db2/V10.1`.
- On Linux operating systems, the default installation directory is `/opt/ibm/db2/V10.1`.

-a *AuthType* (Linux or UNIX)

Represents the authentication type for the instance. *AuthType* can be one of `SERVER`, `CLIENT`, or `SERVER_ENCRYPT`. `SERVER` is the default. This parameter is optional.

-u *FencedID*

Represents the name of the user under which fenced user defined functions (UDFs) and fenced stored procedures will run. This flag is not required if you are creating an instance on a client. Specify the name of the fenced user you created.

InstName

Represents the name of instance. The name of the instance must be the same as the name of the instance owning user. Specify the name of the instance owning user you created. The instance will be created in the instance owning user's home directory.

Example

For example, if you are using server authentication, your fenced user is `db2fenc1`, and your instance owning user is `db2inst1`, use the following command to create an instance on an AIX operating system:

```
/opt/IBM/db2/V10.1/instance/db2icrt -a server -u db2fenc1 db2inst1
```

What to do next

(Optional) After you create an instance you can configure notification for health monitoring. This task can be performed using the DB2 CLP commands.

Operating system user limit requirements (Linux and UNIX)

This topic outlines the recommended operating system user process resource limits (ulimits) on Linux and UNIX operating systems.

Depending on your installation, the DB2 database engine automatically raises the ulimits to varying degrees:

- For root installations, the DB2 database engine automatically raises ulimits where necessary based on the needs of the DB2 database system.
- For non-root installations, the DB2 database engine can update only the **data**, **nofile**, and **fsize** ulimits for the engine process up to the hard limits imposed by the system administrator.

In either case, it might be more practical to set the resource limits permanently on your system. Especially for non-root installations, the **data**, **nofile**, and **fsize** ulimit values should be set appropriately by an administrator after installing.

Recommended ulimit values for non-root installations

After a non-root installation is completed, verify the hard operating system ulimits for the **data**, **nofile**, and **fsize** resources as the instance owner. The recommended values are outlined in the following table:

Table 4. Recommended ulimit values for non-root installations

Hard ulimit resource	Description	Minimum value	Recommended value	Command to query the value
data	Maximum private memory allowed for a process	The amount of memory available on the computer	Unlimited	ulimit -Hd
nofile	Maximum number of open files allowed for a process	Larger than the sum of all MAXFILOP database configuration parameters for all databases in the instance	65536	ulimit -Hn
fsize	Maximum file size allowed	Unlimited	Unlimited	ulimit -Hf

If the minimum ulimit values are not met, the DB2 database engine could encounter unexpected operating system resource shortage errors. These errors can lead to a DB2 database system outage.

Contact a root user or a System Administrator if the ulimit values need to be updated for your non-root installation. Values for ulimits must be set manually, except on AIX operating systems, where you can set ulimits by running the **db2rfe** command.

DB2 database products in a workload partition (AIX)

A DB2 installation process on a system WPAR is similar to any other installation, with the additional benefit of using DB2 database products in a global read-only installation path.

In AIX 6.1 there are two types of workload partitions (WPARs): system WPARs and application WPARs. DB2 installation is only supported on a system WPAR. System WPARs either share the /usr and /opt directories with the global environment, and have a local copy of the /usr and /opt directories.

A DB2 database product can be installed in a local file system on a system WPAR, or a DB2 database product can be installed in a global environment with a DB2 copy shared with other system WPARs. When a DB2 copy is installed in a global environment under either the /usr or /opt directory, which are shared with system WPARs, those system WPARs are able to use the shared DB2 copy to setup DB2 instances.

Each system WPAR manages its own DB2 instances and DAS related to the DB2 copy, and can see only the DB2 instances and DAS created specifically for the system WPAR. DAS and instances created on one WPAR, or in a global environment, are not visible from any other system (system WPAR or global environment).

When a DB2 copy is installed in a global environment, DB2 instances and DAS can be created and managed on each system WPAR sharing the globally installed DB2 copy.

There are specific considerations when a DB2 copy is installed in a global environment:

Installing the DB2 copy

Installing a DB2 copy on a system WPAR is similar to any other DB2 database product installation, with the following exceptions. The following cannot be installed on a system WPAR:

- IBM Tivoli System Automation for Multiplatforms (SA MP)
- IBM Data Studio

Uninstalling the DB2 copy

Before uninstalling the DB2 copy in a global environment, the AIX system WPARs that shares the DB2 copy with DB2 instances or DAS, must be active. In addition, before uninstalling the DB2 copy, any system WPARs sharing the DB2 copy cannot have any related DB2 instances, or DAS associated with, or in use by, that DB2 copy. All instances and DAS must either be:

- dropped (using the **db2idrop** or **dasdrop** command), or,
- updated the instance or DAS to another DB2 copy (using the **db2iupdt** or **dasupdt** command).

Note: The **db2idrop** and **dasdrop** command must be run as the root user.

Applying a fix pack to the DB2 copy

Applying fix packs on a system WPAR is similar to any other DB2 database product update. However, before applying a fix pack to a DB2 copy in a global environment, the AIX system WPARs that share the DB2 copy must be active for the instances to be updated. In addition, before applying a fix pack, any system WPARs sharing the DB2 copy cannot have any related DB2 instances or DAS running. All instances and DAS related to the DB2 copy to be updated must either be:

- stopped (using the **db2stop** command), or,
- run the **installFixPack** command with the **-f update** parameter to bypass all the checking on DB2 library loading, instance and DAS properly stopped or not. (However, this is not recommended.)

Considerations when using the db21s command

When installing the DB2 copy in a global environment, the directory containing the **db21s** command (/usr/local/bin) is linked to a DB2 copy in the global environment. The **db21s** command, is used to list the DB2 Version 9 (or higher) products installed on the system. On a system WPAR, if the /usr directory is shared with the global partition as read-only, the **db21s** command running on the system WPAR might not work unless the link target also exists on the system WPAR, and, the related DB2 copy is registered on the system WPAR. The **db21s** command can be found both in the DB2 installation media and in a DB2 install copy on the system, and can be run from either location.

DB2 database product installation on an NFS-mounted directory

You can install DB2 products on an NFS server and share the installed DB2 products on its NFS client machines. This does not apply to the non-root installation.

If your DB2 product is installed on an NFS-mounted file system, you must verify that Network File System (NFS) is running on each machine.

Setup for shared DB2 on NFS will require a few additional considerations including:

- Licensing (When DB2 product installation is shared by several NFS client systems, manual validation is needed to make sure any systems using the shared DB2 copy have the entitlement).
- Ensuring that the mount point for the DB2 product installation is mounted on the appropriate install path before starting to use the shared DB2.
- Confirming if the link `/usr/local/bin/db2ls` from the related DB2 installation path, which lists installed DB2 products and features, is set up and maintained properly to detect DB2 products and features.
- Performing some manual steps to handle DB2 registries on machines sharing the installed DB2.

For detailed instructions on setting up NFS, see www.ibm.com/developerworks/data/library/long/dm-0609lee.

Note: Starting from DB2 V9.7, the handling of DB2 registries in shared DB2 environment is simplified. In the section C of "Setting Up DB2 Registries" of the white paper mentioned previously, you only need to handle the global profile registry "DB2SYSTEM" on the NFS client machines. The `profiles.reg` (listing the instance names related to the DB2 copy) and `default.env` (storing the global profile registries) are removed from the DB2 installation path. The `profiles.reg` is not needed any more. The global profile registry repository is moved from `default.env` to the `global.reg`. With this change, any steps needed on NFS client systems for `profiles.reg` and `default.env` files in the white paper can be ignored. **The IBM DB2 pureScale Feature does not support shared DB2 installation.**

Installing and setting up OpenSSH

This task describes how to obtain and setup Open Secure Shell (OpenSSH).

Before you begin

Uncomment the entries in the ssh configuration files.

File: `/etc/ssh/ssh_config`

Port 22
Protocol 2

File: `/etc/ssh/sshd_config`

PermitRootLogin yes
PasswordAuthentication no

About this task

You need to perform these steps on each host you want to participate in the DB2 pureScale instance.

For Linux users, OpenSSH is installed by default on SLES 10 SP3 or later, and RHEL 5.5..

Open Secure Shell (OpenSSH) is an open source version of the SSH protocol suite of network connectivity tools. The tools provide shell functions that are authenticated and encrypted. A shell is a command language interpreter that reads

input from a command-line string, stdin, or a file. The steps in this topic will allow you to connect to a remote server through ssh without having to enter a password.

For the DB2 pureScale Feature, you need to set up passwordless SSH access for the root user. The instance owner requires passwordless SSH access, however, the DB2 installation processes set this up if the instance owner does not have it setup. The following steps outline how to set up passwordless SSH access for the root user.

Note: The instance owner must have the rlogin attribute set to the default value of TRUE.

Procedure

1. AIX operating systems only: If OpenSSH is not available on your system, you can obtain it from the latest AIX Expansion Pack and Web Download Pack (<http://www.ibm.com/systems/power/software/aix/expansionpack/index.html>). The OpenSSH fileset includes manual pages with `openssh.man.en_US`. On the Internet, openBSD provides documentation at <http://www.openssh.org/manual.html>.
2. AIX operating systems only: Install OpenSSH. The default setting on AIX for OpenSSH is public key enabled.
3. Setup public key-based authentication. Public key-based authentication enables a single user ID to log in as that same user ID on each host in the instance without being prompted for a password. Public key-based authentication must be enabled for the root ID to use passwordless SSH.

If the user ID has a `~/.ssh` directory, ensure that it does not allow group or other write access. Ensure that the home directory for the user does not allow group or other write access. SSH views this situation as a security exposure and if the directory permissions are not restrictive enough, it does not allow public key-based authentication.

An existing `~/.ssh` directory is not required as the **ssh-keygen** command creates this directory if it does not exist and set up the appropriate access.

From your `~/.ssh` directory, generate a public key/private key pair:

```
$ ssh-keygen -t dsa
```

Whenever prompted for input, hit Enter to accept the default value. (Ensure that no passphrase is entered, or else SSH challenges each authentication attempt, expecting the same passphrase as a response from the user. However, the DB2 product does not allow remote shell utilities to prompt for additional verification.) This action generates two new files in the `~/.ssh` directory, `id_dsa` (the private key) and `id_dsa.pub` (the public key) for DSA encryption.

4. You must generate the public key on each host and append the contents of each public key from each host to a single file called `authorized_keys`. Then copy the `authorized_keys` file to the user's `$HOME/.ssh` directory on each host and run the **chmod 644 authorized_keys** command.

Kernel parameters for DB2 database server installation (Linux and UNIX)

The configuration or modification of kernel parameters for DB2 database server installation depends on your operating system.

Recommended kernel configuration parameters (HP-UX):

For HP-UX systems running a DB2 64-bit database system, run the **db2osconf** command to suggest appropriate kernel configuration parameter values for your system.

The **db2osconf** utility can only be run from *\$DB2DIR/bin*, where *DB2DIR* is the directory where you installed your DB2 database product.

Kernel parameter requirements (Linux):

The database manager uses a formula to automatically adjust kernel parameter settings and eliminate the need for manual updates to these settings.

Interprocess communication kernel parameters

When instances are started, if an interprocess communication (IPC) kernel parameter is below the enforced minimum value, the database manager updates it to enforced minimum value. The IPC kernel parameter values changed when a DB2 instance is started do not persist when the system is rebooted. After a reboot, kernel settings might be lower than the enforced values until a DB2 instance is started. By adjusting any kernel parameter settings, the database manager prevents unnecessary resource errors.

For the latest information about supported Linux distributions, see <http://www.ibm.com/software/data/db2/linux/validate/>.

Table 5. Enforced minimum settings for Linux interprocess communication kernel parameters

IPC kernel parameter	Enforced minimum setting
kernel.shmmni (SHMMNI)	256 * <size of RAM in GB>
kernel.shmmax (SHMMAX)	<size of RAM in bytes> ¹
kernel.shmall (SHMALL)	2 * <size of RAM in the default system page size> ²
kernel.sem (SEMMNI)	256 * <size of RAM in GB>
kernel.sem (SEMMSL)	250
kernel.sem (SEMMNS)	256 000
kernel.sem (SEMOPM)	32
kernel.msgmni (MSGMNI)	1 024 * <size of RAM in GB>
kernel.msgmax (MSGMAX)	65 536
kernel.msgmnb (MSGMNB)	65 536 ³

Table 5. Enforced minimum settings for Linux interprocess communication kernel parameters (continued)

IPC kernel parameter	Enforced minimum setting
1. On 32-bit Linux operating systems, the enforced minimum setting for SHMMAX is limited to 4 294 967 295 bytes.	
2. SHMALL limits the total amount of virtual shared memory that can be allocated on a system. Each DB2 data server efficiently manages the amount of system memory it consumes, also known as committed memory. The DB2 data server allocates more virtual memory than it commits to support memory preallocation and dynamic memory management. Memory preallocation benefits performance. Dynamic memory management is the process of growing and shrinking real memory usage within separate virtual shared memory areas. To support memory preallocation and dynamic memory management effectively data servers frequently have to allocate more virtual shared memory on a system than the amount of physical RAM. The kernel requires this value as a number of pages.	
3. Load performance might benefit from a larger message queue size limit, specified in bytes by MSGMNB . You can view message queue usage by running the ipcs -q command. If the message queues are at capacity, or reaching capacity, during load operations, consider increasing the number of bytes the message queue size limit.	

Other recommended kernel parameter settings

Other recommended kernel parameter settings are listed in the following table.

Table 6. Configuring other Linux kernel parameters

Recommended kernel parameter setting	Configuring the kernel parameters for DB2 data server
vm.swappiness=0	This parameter defines how prone the kernel is to swapping application memory out of physical random access memory (RAM). The default setting is vm.swappiness=60 . The recommended kernel parameter setting, vm.swappiness=0 , configures the kernel to give preference to keeping application memory in RAM instead of assigning more memory for file caching. This setting avoids unnecessary paging and excessive use of swap space. This setting is especially important for data servers configured to use the self-tuning memory manager (STMM).

Table 6. Configuring other Linux kernel parameters (continued)

Recommended kernel parameter setting	Configuring the kernel parameters for DB2 data server
vm.overcommit_memory=0	This parameter influences how much virtual memory the kernel permits allocating. The default setting, vm.overcommit_memory=0 , sets the kernel to disallow individual processes from making excessively large allocations, however the total allocated virtual memory is unlimited. Having unlimited virtual memory is important for DB2 data servers, which retain additional unused virtual memory allocations for dynamic memory management. Unreferenced allocated memory is not backed by RAM or paging space on Linux systems. Avoid setting vm.overcommit_memory=2 , as this setting limits the total amount of virtual memory that can be allocated, which can result in unexpected errors.

Modifying kernel parameters (HP-UX):

For your DB2 database product to perform properly on HP-UX, you might need to update your system's kernel configuration parameters. If you update your kernel configuration parameter values, you must restart your computer.

Before you begin

You must have root user authority to modify kernel parameters.

Procedure

To modify kernel parameters:

1. Enter the **sam** command to start the System Administration Manager (SAM) program.
2. Double-click the **Kernel Configuration** icon.
3. Double-click the **Configurable Parameters** icon.
4. Double-click the parameter that you want to change and type the new value in the **Formula/Value** field.
5. Click **OK**.
6. Repeat these steps for all of the kernel configuration parameters that you want to change.
7. When you are finished setting all of the kernel configuration parameters, select **Action > Process New Kernel** from the action menu bar.

Results

The HP-UX operating system automatically restarts after you change the values for the kernel configuration parameters.

Tip:

kctune can also be used on HP-UX for adjusting kernel parameters.

Modifying kernel parameters (Linux):

The database manager uses a formula to automatically adjust kernel parameter settings and eliminate the need for manual updates to these settings.

Before you begin

You must have root authority to modify kernel parameters.

Procedure

To update kernel parameters on Red Hat and SUSE Linux:

1. Run the **ipcs -l** command to list the current kernel parameter settings.
2. Analyze the command output to determine whether you have to change kernel settings or not by comparing the current values with the *enforced minimum settings* at <http://publib.boulder.ibm.com/infocenter/db2luw/v9r8/topic/com.ibm.db2.luw.qb.server.doc/doc/c0057140.html>. The following text is an example of the **ipcs** command output with comments added after **//** to show what the parameter names are:

```
# ipcs -l

----- Shared Memory Limits -----
max number of segments = 4096           // SHMMNI
max seg size (kbytes) = 32768           // SHMMAX
max total shared memory (kbytes) = 8388608 // SHMALL
min seg size (bytes) = 1

----- Semaphore Limits -----
max number of arrays = 1024             // SEMMNI
max semaphores per array = 250          // SEMMSL
max semaphores system wide = 256000    // SEMMNS
max ops per semop call = 32             // SEMOPM
semaphore max value = 32767

----- Messages: Limits -----
max queues system wide = 1024          // MSGMNI
max size of message (bytes) = 65536    // MSGMAX
default max size of queue (bytes) = 65536 // MSGMNB
```

- Beginning with the first section on Shared Memory Limits, the **SHMMAX** limit is the maximum size of a shared memory segment on a Linux system. The **SHMALL** limit is the maximum allocation of shared memory pages on a system.
 - It is recommended to set the **SHMMAX** value to be equal to the amount of physical memory on your system. However, the minimum required on x86 systems is 268435456 (256 MB) and for 64-bit systems, it is 1073741824 (1 GB).
 - The **SHMALL** parameter is set to 8 GB by default (8388608 KB = 8 GB). If you have more physical memory than 8 GB, and it is to be used for DB2, then this parameter increases to approximately 90% of your computer's physical memory. For instance, if you have a computer system with 16 GB of memory to be used primarily for DB2, then **SHMALL** should be set to 3774873 (90% of 16 GB is 14.4 GB; 14.4 GB is then divided by 4 KB, which is the base page size). The **ipcs** output converted **SHMALL** into kilobytes. The kernel requires this value as a number of pages. If you are upgrading to DB2 Version 10.1 and you are not using the default **SHMALL** setting, you must increase the **SHMALL** setting by an additional 4 GB. This increase in memory is required by the fast communication manager (FCM) for additional buffers or channels.

- The next section covers the amount of semaphores available to the operating system. The kernel parameter **sem** consists of four tokens, **SEMMSL**, **SEMMNS**, **SEMOPM** and **SEMMNI**. **SEMMNS** is the result of **SEMMSL** multiplied by **SEMMNI**. The database manager requires that the number of arrays (**SEMMNI**) be increased as necessary. Typically, **SEMMNI** should be twice the maximum number of agents expected on the system multiplied by the number of logical partitions on the database server computer plus the number of local application connections on the database server computer.
 - The third section covers messages on the system.
 - The **MSGMNI** parameter affects the number of agents that can be started; the **MSGMAX** parameter affects the size of the message that can be sent in a queue, and the **MSGMNB** parameter affects the size of the queue.
 - The **MSGMAX** parameter should be changed to 64 KB (that is, 65536 bytes), and the **MSGMNB** parameter should be increased to 65536.
3. Modify the kernel parameters that you have to adjust by editing the `/etc/sysctl.conf` file. If this file does not exist, create it. The following lines are examples of what should be placed into the file:
- ```
#Example for a computer with 16GB of RAM:
kernel.shmmni=4096
kernel.shmmax=17179869184
kernel.shmall=8388608
#kernel.sem=<SEMMSL> <SEMMNS> <SEMOPM> <SEMMNI>
kernel.sem=250 1024000 32 4096
kernel.msgmni=16384
kernel.msgmax=65536
kernel.msgmnb=65536
```
4. Run **sysctl** with **-p** parameter to load in **sysctl** settings from the default file `/etc/sysctl.conf`:
- ```
sysctl -p
```
5. Optional: Have the changes persist after every reboot:
- (SUSE Linux) Make `boot.sysctl` active.
 - (Red Hat) The `rc.sysinit` initialization script reads the `/etc/sysctl.conf` file automatically.
- For the latest information about supported Linux distributions, see <http://www.ibm.com/software/data/db2/linux/validate/>.

Modifying kernel parameters (Solaris):

For the DB2 database system to operate properly, it is recommended that you update your system's kernel configuration parameters. You can use the **db2osconf** utility to suggest recommended kernel parameters. If you want to take advantage of project resource controls (`/etc/project`), consult your Solaris documentation.

Before you begin

You must have root authority to modify kernel parameters.

To use the **db2osconf** command, you must first install the DB2 database system. The **db2osconf** utility can only be run from `$DB2DIR/bin`, where `DB2DIR` is the directory where you installed your DB2 database product.

You must restart your system after modifying kernel parameters.

Procedure

To set a kernel parameter:

Add a line at the end of the `/etc/system` file as follows:

```
set parameter_name = value
```

For example, to set the value of the `msgsys:msginfo_msgmax` parameter, add the following line to the end of the `/etc/system` file:

```
set msgsys:msginfo_msgmax = 65535
```

What to do next

After updating the `/etc/system` file, restart the system.

Prerequisites for a DB2 database server installation (Windows)

Before you install your DB2 database product, ensure that the prerequisites for each of the tasks are met for installing on Windows operating systems.

Required user accounts for installation of DB2 server products (Windows)

DB2 server products have installation prerequisites regarding user accounts that you must consider before you begin installation tasks.

- An installation user account and
- Optional - one or more setup user accounts. You can create these accounts during the installation.
 - A DB2 Administration Server (DAS) user account
 - A DB2 instance user account. You can also use the LocalSystem account for products other than DB2 Enterprise Server Edition.

The installation user account is the account of the user performing the installation. The installation user account must be defined before running the DB2 Setup wizard. The setup user accounts can be defined before installation or you can have the DB2 Setup wizard create them for you.

All user account names must adhere to your system naming rules and to DB2 User, user ID and group naming rules.

If you use an installation user account that contains non-English characters which are not specified in DB2 naming rules, the DB2 installation will fail.

Extended security on Windows

DB2 database products offer extended Windows security. If the extended security feature is selected, you must add the users who will administer or use the DB2 database product to either the DB2ADMNS or DB2USERS group as appropriate.

The DB2 installer creates these two new groups. You can either specify a new name or accept the default names during installation.

To enable this security feature, select the **Enable operating system security** check box on the **Enable operating system security for DB2 objects** panel during the DB2 installation. Accept the default values for the DB2 Administrators Group field, and the DB2 Users Group field. The default group names are DB2ADMNS and

DB2USERS. If there is a conflict with existing group names, you will be prompted to change the group names. If required, you can specify your own group names.

DB2 server user accounts

Installation user account

A local or domain user account is required to perform the installation. Normally, the user account must belong to the *Administrators* group on the computer where you will perform the installation.

Alternatively, a non-Administrator user account can be used. This alternative requires that a member of the Windows Administrators group first configure the Windows elevated privileges settings to allow a non-Administrator user account to perform an installation.

On Windows 2008 and Windows Vista or higher, a non-administrator can perform an installation, but will be prompted for administrative credentials by the DB2 Setup wizard.

The user right "Access this computer from the network" is required for the installation user account.

The installation user ID must belong to the Domain Administrators group on the domain if the installation requires a domain account to be created or verified.

You may also use the built-in LocalSystem account as your Service Logon account for all products, except DB2 Enterprise Server Edition.

User rights granted by the DB2 installer

The DB2 installation program does not grant the Debug Programs user right. The DB2 installer grants the following user rights:

- Act as part of the operating system
- Create token object
- Lock pages in memory
- Log on as a service
- Increase quotas
- Replace a process level token

DB2 Administration Server (DAS) user account

A local or domain user account is required for the DB2 Administration Server (DAS).

Important: The DB2 Administration Server (DAS) has been deprecated in Version 9.7 and might be removed in a future release. The DAS is not supported in DB2 pureScale environments. Use software programs that use the Secure Shell protocol for remote administration. For more information, see "DB2 administration server (DAS) has been deprecated" at .

If you are performing a response file installation, you can also specify the Local System account in the response file. For more details, refer to the sample response files in the db2\windows\samples directory.

The LocalSystem account is available for all products, except DB2 Enterprise Server Edition and can be selected through the DB2 Setup wizard.

The DAS is a special DB2 administration service used to support the GUI tools and assist with administration tasks on local and remote DB2 servers.

The DAS has an assigned user account that is used to log the DAS service on to the computer when the DAS service is started.

You can create the DAS user account before installing DB2 or you can have the DB2 Setup wizard create it for you. If you want to have the DB2 Setup wizard create a new domain user account, the user account you use to perform the installation must have authority to create domain user accounts. The user account must belong to the *Administrators* group on the computer where you will perform the installation. This account will be granted the following user rights:

- Act as part of the operating system
- Debug programs
- Create token object
- Lock pages in memory
- Log on as a service
- Increase quotas (adjust memory quotas for a process on Windows XP and Windows Server 2003 operating systems)
- Replace a process level token

If extended security is enabled, the DB2ADMNS group will have all these privileges. You can add users to that group and you do not need to add these privileges explicitly. However, the user still needs to be a member of the Local Administrators group.

The "Debug programs" privilege is only needed when DB2 group lookup is explicitly specified to use the access token.

If the user account is created by the install program, the user account will be granted these privileges and if the user account already exists, this account will also be granted these privileges. If the install grants the privileges, some of them will only be effective on first log on by the account that was granted the privileges or upon reboot.

It is recommended that the DAS user have SYSADM authority on each of the DB2 database systems within your environment so that it can start or stop other instances if required. By default, any user that is part of the *Administrators* group has SYSADM authority.

DB2 instance user account

The user account must belong to the *Administrators* group on the computer where you will perform the installation.

A local or domain user account is required for the DB2 instance because the instance is run as a Windows service and the service will be executing in the security context of the user account. When you use a domain user account to perform a database operation (such as, creating a database) against a DB2 instance, the DB2 service needs to access the domain to authenticate and search for the user's group membership. By default, a domain will only allow a domain user to query the domain and hence, the DB2 service needs to be running in the security context of a domain user. An error will occur if you use a domain user account to perform a database operation against a DB2 service running with either a Local user account or a LocalSystem account.

You may also use the built-in LocalSystem account to run the installation for all products, except for DB2 Enterprise Server Edition.

You can create the DB2 instance user account before installing DB2 or you can have the DB2 Setup wizard create it for you. If you want to have the

DB2 Setup wizard create a new domain user account, the user account you use to perform the installation must have authority to create domain user accounts. This account will be granted the following user rights:

- Act as part of the operating system
- Debug programs
- Create token object
- Increase quotas
- Lock pages in memory
- Log on as a service
- Replace a process level token

If extended security is enabled, then the DB2ADMNS group will have all these privileges. You can add users to that group and you do not need to add these privileges explicitly. However, the user still needs to be a member of the Local Administrators group.

The "Debug programs" privilege is only needed when DB2 group lookup is explicitly specified to use the access token.

If the user account is created by the install program, the user account will be granted these privileges and if the user account already exists, this account will also be granted these privileges. If the install grants the privileges, some of them will only be effective on first log on by the account that was granted the privileges or upon reboot.

DB2 services running on your system (Windows)

During DB2 product installation, you might find it beneficial to be aware of all services that are running and which services are needed by the process.

The following table lists the DB2 services that run on your system when a DB2 product is installed:

Table 7. DB2 Services

Service Display Name	Service Name	Description
DB2 - (DB2 copy name) - <instance name> [<-nodenum>] where <-nodenum> is added for DB2 ESE instances.	<instance name>[<-nodenum>]	Allows applications to create, update, control, and manage DB2 databases.
DB2 Governor (DB2 copy name)	DB2GOVERNOR _db2copyname	Collects statistics for applications connected to DB2 databases.
DB2 Information Center Server	DB2ICSERVER_Vxx (where xx is the DB2 version you are running, for example, V10 is Version 10.1)	Provides documentation for DB2 products.
DB2 License Server (DB2 copy name)	DB2LICD _db2copyname	Monitors DB2 license compliance.
DB2 Management Service (DB2 copy name)	DB2MGMTSVC _db2copyname	Manages DB2 registry entries for compatibility with earlier DB2 copy versions.

Table 7. DB2 Services (continued)

Service Display Name	Service Name	Description
DB2 Remote Command Server (DB2 copy name)	DB2REMOTECMD_db2copyname	Supports remote DB2 command execution.
DB2DAS - DB2DASXX	DB2DASXX (where XX is 00-99)	Supports local and remote database administrative requests.
DB2TS - (DB2 copy name) - <instance name> [<-nodenum>]	<instance name> [<-nodenum>]	Indexes and searches text documents in DB2 databases.
DB2EXT - (DB2 copy name) - <instance name> [<-nodenum>]	<instance name> [<-nodenum>]	
IBM Secure Shell Server for Windows	ibmmsshd	IBM Secure Shell Server for Windows

Restricting operating system privileges of the db2fmp process (Windows)

On Windows operating systems, you can restrict privileges of the **db2fmp** process to the privileges assigned to the DB2USERS group.

About this task

Restrictions

This option is not available if LocalSystem is selected as the service account.

Procedure

- On Windows operating systems, if you have Extended Security enabled, to restrict the **db2fmp** process to the privileges assigned to the DB2USERS group:
 - Run the **db2set** command and set **DB2_LIMIT_FENCED_GROUP** to ON. By default, this registry variable is set to OFF.

```
db2set DB2_LIMIT_FENCED_GROUP = ON
```
 - Add the DB2 service account into the DB2USERS group.
- You can grant additional operating system privileges to the **db2fmp** process by following these steps:
 - Create a new user group, or choose an existing user group (for example, db2FencedGroup).
 - Add the DB2 service account into the group.

Results

In addition to the privilege of DB2USERS group, the **db2fmp** process has the operating system privilege of the chosen user group chosen.

Setting up Windows elevated privileges before installing a DB2 product (Windows)

The usual method to install a DB2 database product on Windows is to use an Administrator user account. However, DB2 database products can be installed using a non-administrator account. To do so, a Windows Administrator must configure the elevated privileges feature in Windows.

About this task

This task explains how a Windows Administrator can set up a computer with elevated privileges to allow installation using a non-Administrator user account. The related task of granting DB2 administration authorities to non-Administrator users is also covered.

Typically a Windows Administrator would perform this task to enable another person who does not have an Administrator account to install a DB2 database product. The role of this person might be only to install DB2 database products or to also administer DB2 database products once installed.

Restrictions

Before initiating this procedure, note the following restrictions on non-Administrator installation using elevated privileges:

- Non-Administrator users can only install fix packs, add-on products, or upgrade DB2 database products if prior installations or upgrades were also performed by the same non-Administrator user.
- Non-Administrator users cannot uninstall a DB2 database product. Those non-Administrator users on a Windows Vista (and later) operating system **can** uninstall a DB2 database product.

This procedure uses the Windows Group Policy Editor.

Procedure

1. Click **Start > Run** and type **gpedit.msc**. The Group Policy window opens.
2. Click on **Computer Configuration > Administrative Templates > Windows Components > Windows Installer**.
3. Enable the following Group Policy settings:
 - Always install with elevated privileges (mandatory)
 - Enable user control over installs (mandatory)
 - Disable Windows Installer. Then set it to Never.
 - Enable user to patch elevated products (optional)
 - Enable user to use media source while elevated (optional)
 - Enable user to browse for source while elevated (optional for new installations, mandatory for fix pack upgrades)
4. Enable elevated privileges for the user account that will be performing the installation.
 - a. Click **User Configuration > Administrative Templates > Windows Components > Windows Installer**.
 - b. Enable the **Always install with elevated privileges (mandatory)** Group Policy setting.
5. Perform setup related to the user account that will install the DB2 database product.
 - Identify the user account that will install the DB2 database product. If necessary, create that account.
 - Give that account write permission for the drive on which an installation is planned.
6. Optional: Complete additional steps applicable to installing fix packs:
 - a. Provide read access to the `sql1ib\cfg` directory.

- b. Ensure that `allowlockdownpatch` is enabled (as described in the Windows Installer SDK documentation) because fix pack installations are considered minor upgrades to the product.
7. Refresh the computer's security policy in any one of the following ways:
 - Reboot the PC.
 - At the command line, enter **gpupdate.exe**.

Results

By following this procedure you will have set up the computer with elevated privileges and set up a user account that will be able to install DB2 database server products, clients and fix packs.

After DB2 database product installation is complete:

- Any user in the system administrative (SYSADM) or system control (SYSCTRL) authority group defined in the database manager configuration for the instance can create and use DB2 databases within the DB2 instance.
- Only a user with local Administrator authority can run DB2 instance utilities, such as **db2icrt**, **db2idrop**, **db2iupdt**, or **db2iupgrade**.
- The authorization requirements for running the **db2start** or **db2stop** command is defined in the topics **START DATABASE MANAGER** command, and **STOP DATABASE MANAGER** command.

What to do next

Using regedit instead of the Windows Group Policy Editor

An alternative to using the Windows Group Policy Editor is to use **regedit**.

1. In the registry branch `HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows`, add the key `installer`
2. Edit the key `installer` with the following values:
 - For `AlwaysInstallElevated`, enter `REG_DWORD=1`
 - For `AllowLockdownBrowse`, enter `REG_DWORD=1`
 - For `AllowLockdownMedia`, enter `REG_DWORD=1`
 - For `AllowLockdownPatch`, enter `REG_DWORD=1`
 - For `DisableMSI`, enter `REG_DWORD=0`
 - For `EnableUserControl`, enter `REG_DWORD=1`
3. In the registry branch `HKEY_CURRENT_USER\SOFTWARE\Policies\Microsoft\Windows`, add the key `installer`
4. Edit the key `installer` with the following values:
 - For `AlwaysInstallElevated`, enter `REG_DWORD=1`

Removing elevated privileges

After you have given elevated privileges, you can reverse this action. To do so, remove the registry key `Installer` under `HKEY_LOCAL_MACHINE\SOFTWARE\Policies\Microsoft\Windows`.

Granting a non-administrator user DB2 administration authorities

At this point, only members of the Windows Administrators group will have DB2 administration authorities. The Windows Administrator has the

option to grant one or more DB2 authorities, such as SYSADM, SYSMAINT, or SYSCTRL to the non-Administrator user who installed the DB2 database product.

Granting user rights (Windows)

This topic describes the steps required to grant user rights on Windows operating systems. Specific user rights are recommended for user accounts required to install and set up DB2.

About this task

To grant advanced user rights on Windows you must be logged on as a local Administrator.

Procedure

1. Click **Start** -> **Run** and type `secpol.msc`. On Windows 2008 and Windows Vista (or higher), click **Start** and type `secpol.msc` in the search bar. Click **OK**.
2. Select **Local Security Policy**.
3. In the left window pane, expand the **Local Policies** object, then select **User Rights Assignment**.
4. In the right window pane, select the user right that you want to assign.
5. From the menu, select **Action** -> **Security...**
6. Click **Add**, then select a user or group to assign the right to, and click **Add**.
7. Click **OK**.

What to do next

If your computer belongs to a Windows domain, the domain user rights can override your local settings. In this case, your Network Administrator must make the changes to the user rights.

Extending the Active Directory Schema for LDAP directory services (Windows)

If you plan to use the Lightweight Directory Access Protocol (LDAP) directory server feature with Windows Server 2003, you have to extend the Active Directory schema to contain DB2 object classes and attribute definitions using the **db2schex** command.

About this task

Extending the directory schema before installing DB2 database products and creating databases provide the following benefits:

- The default DB2 instance, created during the installation, is cataloged as a DB2 node in Active Directory, provided that the installation user ID had sufficient privileges to write to Active Directory.
- Any databases created after installation is automatically cataloged into Active Directory.

Procedure

To extend the directory schema:

1. Log onto any machine that is part of the Windows domain with a Windows user account that has Schema Administration authority.

- Run the **db2schex** command from the installation DVD . You can run this command without logging off and logging on again, as follows:

```
runas /user:MyDomain\Administrator x:\db2\Windows\utilities\db2schex.exe
```

where x: represents the DVD drive letter.

What to do next

When **db2schex** completes, you can proceed with the installation of your DB2 database product; or if you have already installed DB2 database products or created databases, you have to manually register the node and catalog the databases. For more information, see the “Enabling LDAP support after DB2 installation is complete” topic.

Installation requirements for DB2 servers and IBM data server clients (AIX)

Before you install DB2 database products on AIX operating systems, ensure that the system you choose meets the necessary operating system, hardware, software, and communications requirements. The **db2prereqcheck** command checks whether your system meets the prerequisites for the installation.

DB2 database products are available in several different editions. Some DB2 database products and features are only available on specific operating systems.

If you are planning to install the IBM DB2 pureScale Feature, different installation prerequisites are required. See the Installing the DB2 pureScale Feature topic.

To install a DB2 database product, the following requirements must be met:

Table 8. AIX installation requirements

Operating System	Hardware
<p>AIX Version 6.1²</p> <ul style="list-style-type: none"> 64-bit AIX kernel is required AIX 6.1 Technology Level (TL) 6 and Service Pack (SP) 5 Minimum C++ runtime level requires the xlc.rte 11.1.0.1 and xlc AIX rte 11.1.0.1 (or later) filesets. <p>AIX Version 7.1</p> <ul style="list-style-type: none"> 64-bit AIX kernel is required AIX 7.1 Technology Level (TL) 0 and Service Pack (SP) 3 Minimum C++ runtime level requires the xlc.rte 11.1.0.1 and xlc AIX rte 11.1.0.1 (or later) filesets. 	<p>64-bit Common Hardware Reference Platform (CHRP) architecture, excluding POWER3 processor-based systems.¹</p> <p>All processors that are capable of running the supported AIX operating systems.</p>

- ¹ To verify that it is a CHRP architecture system, issue the command **lscfg** and look for the following output: Model Architecture: chrp. For POWER3 processor-based systems, first upgrade to POWER4 processor-based systems. POWER3 processor-based systems are not supported.
- ² In AIX 6.1 there are two types of Workload Partitions (WPARs): system WPARs and application WPARs. DB2 installation is supported only on a system WPAR. AIX 6.1 also supports the ability to encrypt a JFS2 file system or set of files.

- ³ Additional considerations for POWER7[®] systems:
 - DB2 products running on POWER7 systems support Active Memory[™] Expansion.
- 32-bit clients are also supported

Note: If the minimum operating system requirements were applied using an operating system upgrade rather than using a new operating system installation, you must install I/O completion ports (IOCP) separately. You can obtain the IOCP libraries from the AIX installation CD and configure IOCP on the DB2 server. For a client only installation, IOCP is not required. For more information, see “configuring IOCP on AIX” in *Troubleshooting and Tuning Database Performance*.

Software considerations

- For application development and runtime considerations, see the topics in Supported programming languages and compilers for database application development.
- You can download the latest IBM C++ Runtime Environment Components for AIX at the IBM AIX XL C and C++ support website.
- (Clients only) If you plan to use Kerberos Authentication, you require IBM Network Authentication Service client v1.4 or later.
- Use the **bosboot** command to switch to the 64-bit kernel.

To switch to a 64-bit kernel, you require root authority and enter the following commands:

```
ln -sf /usr/lib/boot/unix_64 /unix
ln -sf /usr/lib/boot/unix_64 /usr/lib/boot/unix
bosboot -a
shutdown -Fr
```

- One of the following browsers is required to view online help and to run First Steps (**db2fs**):
 - Firefox 3.0 and later
 - Google Chrome
 - Safari 4.0
- An X Window System software capable of rendering a graphical user interface is required if:
 - you want to use the DB2 Setup wizard to install a DB2 database product on Linux or UNIX operating systems
- For details regarding known AIX issues, see www.ibm.com/support/docview.wss?uid=swg21165448

Mounting considerations

For root installations of DB2 products on UNIX operating systems, do not mount your file system with the *noatime* option.

Installation requirements for DB2 servers and IBM data server clients (HP-UX)

Before you install DB2 database products on HP-UX operating systems, ensure that the system you choose meets the necessary operating system, hardware, software, and communications requirements. The **db2prereqcheck** command checks whether your system meets the prerequisites for the installation.

DB2 database products are available in several different editions. Some DB2 database products and features are only available on specific operating systems.

To install a DB2 product on an HP-UX operating system, the following operating system, hardware, and communications requirements must be met:

Note: A 64-bit HP-UX operating system is required to support DB2 products.

Table 9. HP-UX installation requirements

Operating System	Hardware
DB2 products are supported on: <ul style="list-style-type: none">• HP-UX 11iv3 (11.31) with:<ul style="list-style-type: none">– PHSS_37202– PHKL_41481– PHKL_42035– PHKL_42335– PHKL_41588– PHSS_41496	Itanium based HP Integrity Series Systems Any hardware that is explicitly compatible and fully capable of running the specified operating system, all the corresponding supporting software, and any associated applications unmodified.

Kernel configuration considerations

A system restart is required if you update the kernel configuration parameters. The kernel configuration parameters are set in `/etc/system`. Depending on the values of your kernel configuration parameters, modify some of them before you install the client or DB2 server products. If the kernel parameter being modified is not listed as dynamic, a system reboot is required to make the changes to `/etc/system` take effect.

Software considerations

- `libpam.so.0` (32-bit) is required for DB2 database servers to run 32-bit non-SQL routines.
- (Clients only) If you plan to use Kerberos authentication, your DB2 applications must be 64-bit.
- (Client and server) To use Kerberos authentication, install the HP `krb5` Kerberos support via the Kerberos packages on the operating system.
- One of the following browsers is required to view online help and to run First Steps (**db2fs**):
 - Firefox 3.0 and later
 - Google Chrome
 - Safari 4.0
- An X Window System software capable of rendering a graphical user interface is required if you want to use the DB2 Setup wizard to install a DB2 product.
- For details regarding known HP-UX issues, see www.ibm.com/support/docview.wss?uid=swg21257602

Note: DB2 products installed on the HP-UX operating system support long host names. The length has been extended to 255 bytes, in any combination of characters or digits.

To enable long host name support, complete the following tasks:

1. Turn on the kernel tunable parameter `expanded_node_host_name`.
`Kctune expanded_node_host_name=1`
2. Compile applications requiring long host name support with the `-D_HPUX_API_LEVEL=20040821` option.

Mounting considerations

For root installations of DB2 products on UNIX operating systems, do not mount your file system with the *noatime* option.

Installation requirements for DB2 servers and IBM data server clients (Linux)

Before you install DB2 database products on Linux operating systems, ensure that the system you choose meets the necessary operating system, hardware, software, and communications requirements. The **db2prereqcheck** command checks whether your system meets the prerequisites for the installation.

DB2 database products are available in several different editions. Some DB2 database products and features are only available on specific operating systems.

If you are planning to install the IBM DB2 pureScale Feature, different installation prerequisites are required. See the Installing the DB2 pureScale Feature topic.

DB2 database products are supported on the following hardware:

- x86 (Intel Pentium, Intel Xeon, and AMD) 32-bit Intel and AMD processors
- x64 (64-bit AMD64 and Intel EM64T processors)
- POWER (System i®, System p®, and POWER Systems that support Linux)
- System z®: zEnterprise (z196, z114, zBX), System z10® or System z9®

The minimum supported operating systems for Linux include:

- Red Hat Enterprise Linux (RHEL) 5 Update 6
- Red Hat Enterprise Linux (RHEL) 6
- SUSE Linux Enterprise Server (SLES) 10 Service Pack 3
- SUSE Linux Enterprise Server (SLES) 11 Service Pack 1
- Ubuntu 10.04

For the latest information about supported Linux distributions, see <http://www.ibm.com/software/data/db2/linux/validate/>.

Note:

- Version 9.7 Fix Pack 2 and later fix packs (and Version 9.8 Fix Pack 2 and later fix packs) include an updated version of IBM Tivoli System Automation for Multiplatforms (SA MP) Base Component that you can use in environments with SLES 11 or POWER7 systems. For more information, see “Installing IBM Tivoli System Automation for Multiplatforms (SA MP) Base Component” or “Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP) Base Component”.

Multithreaded architecture limitations

If you are installing a DB2 32-bit database product on a Linux operating system, consider upgrading to a 64-bit operating system and installing the DB2 64-bit

database product instead. The multithreaded architecture generally simplifies memory configuration. However, this could affect the memory configuration of 32-bit DB2 database servers. For example:

- Private memory for agent threads is allocated within a single process. The aggregate of all private memory allocations for database agents might not fit in a single process memory space.
- Support for multiple databases is limited because all database shared memory segments for all databases are allocated in a single process. You might need to reduce the memory usage for some databases in order to activate all databases successfully at the same time. However, the database manager performance might be impacted. Another alternative is to create multiple instances and catalog the databases across the instances. However, sufficient system resources is required to support this configuration.

Distribution Requirements

You should update your kernel configuration parameters before installing a DB2 database product. The default values for particular kernel parameters might not be sufficient when running a DB2 database system.

You might also have other products or applications that require Linux system resources. You should modify the kernel configuration parameters based on the needs of your Linux system working environment.

The kernel configuration parameters are set in `/etc/sysctl.conf`.

Refer to your operating system manual for information about setting and activating these parameters using the **sysctl** command.

Package requirements

The following tables list the package requirements for SLES and RHEL distributions:

- `libpam.so.0` (32-bit) is required for DB2 database servers to run 32-bit non-SQL routines.
- `libaio.so.1` is required for DB2 database servers using asynchronous I/O.
- `libstdc++.so.6` is required for DB2 database servers and clients.

Package requirements for SLES and RHEL

Package name	Description
<code>libaio</code>	Contains the asynchronous library required for DB2 database servers.
<code>compat-libstdc++</code>	Contains <code>libstdc++.so.6</code> (not required for Linux on POWER)

The following tables list the package requirements for SUSE Linux and Red Hat distributions for DB2 partitioned database servers.

- The **ksh93** Korn shell is required for SUSE10 and RHEL5 systems. The **pdksh** Korn Shell package is required for all other DB2 database systems.
- A remote shell utility is required for partitioned database systems. DB2 database systems support the following remote shell utilities:
 - **rsh**
 - **ssh**

By default, DB2 database systems use **rsh** when executing commands on remote DB2 nodes, for example, when starting a remote DB2 database partition. To use the DB2 database system default, the **rsh-server** package must be installed (see following table). More information about **rsh** and **ssh** is available in the DB2 Information Center.

If you choose to use the **rsh** remote shell utility, **inetd** (or **xinetd**) must be installed and running as well. If you choose to use the **ssh** remote shell utility, you need to set the **DB2RSHCMD** communication variable immediately after the DB2 installation is complete. If this registry variable is not set, **rsh** is used.

- The **nfs-utils** Network File System support package is required for partitioned database systems.

All required packages should be installed and configured before continuing with the DB2 database system setup. For general Linux information, see your Linux distribution documentation.

Package requirements for SUSE Linux

Package name	Description
pdksh or ksh93	Korn Shell.
openssh	This package contains a set of server programs which allow users to run commands on (and from) remote computers via a secure shell. This package is not required if you use the default configuration of DB2 database systems with rsh .
rsh-server	This package contains a set of server programs which allow users to run commands on remote computers, login in to other computers, and copy files between computers (rsh , rexec , rlogin , and rcp). This package is not required if you configure DB2 database systems to use ssh .
nfs-utils	Network File System support package. It allows access to local files from remote computers.

Package requirements for Red Hat

Directory	Package name	Description
/System Environment/Shell	pdksh or ksh93	Korn Shell.
/Applications/Internet	openssh	This package contains a set of client programs which allow users to run commands on a remote computer via a Secure Shell. This package is not required if you use the default configuration of DB2 database systems with rsh .
/System Environment/Daemons	openssh-server	This package contains a set of server programs which allow users to run commands from a remote computer via a Secure Shell. This package is not required if you use the default configuration of DB2 database systems with rsh .
/System Environment/Daemons	rsh-server	This package contains a set of programs which allow users to run commands on a remote computer. Required for partitioned database environments. This package is not required if you configure DB2 database systems to use ssh .

Package requirements for Red Hat

Directory	Package name	Description
/System Environment/Daemons	nfs-utils	Network File System support package. It allows access to local files from remote computers.

Software considerations

- (Client and Server) To use Kerberos authentication, install the Linux krb5 Kerberos support via the operating system Kerberos packages.
- One of the following browsers is required to view online help and to run First Steps (**db2fs**):
 - Firefox 3.0 and later
 - Google Chrome
 - Safari 4.0
- An X Window System software capable of rendering a graphical user interface is required if:
 - you want to use the DB2 Setup wizard to install a DB2 database product on Linux or UNIX operating systems, or
 - you want to use any DB2 graphical tools on Linux for x86 and Linux on AMD 64/EM64T.
- Micro Focus does not offer support for any of its COBOL compiler products on SLES 11.

Security-enhanced Linux considerations

On RHEL systems, if Security-enhanced Linux (SELinux) is enabled and in enforcing mode, the installer might fail due to SELinux restrictions.

To determine if SELinux is installed and in enforcing mode, you can do one of the following:

- check the `/etc/sysconfig/selinux` file
- run the **sestatus** command
- check the `/var/log/messages` file for SELinux notices.

To disable SELinux, you can do one of the following:

- set it in permissive mode and run the **setenforce 0** command as a superuser
- modify `/etc/sysconfig/selinux` and reboot the machine.

If your DB2 database product installs successfully on a RHEL system, DB2 processes will run in the unconfined domain. To assign DB2 processes to their own domains, modify the policy. A sample SELinux policy is provided in the `sqlib/samples` directory.

Modifying kernel parameters (Linux)

The database manager uses a formula to automatically adjust kernel parameter settings and eliminate the need for manual updates to these settings.

Before you begin

You must have root authority to modify kernel parameters.

Procedure

To update kernel parameters on Red Hat and SUSE Linux:

1. Run the **ipcs -l** command to list the current kernel parameter settings.
2. Analyze the command output to determine whether you have to change kernel settings or not by comparing the current values with the *enforced minimum settings* at <http://publib.boulder.ibm.com/infocenter/db2luw/v9r8/topic/com.ibm.db2.luw.qb.server.doc/doc/c0057140.html>. The following text is an example of the **ipcs** command output with comments added after **//** to show what the parameter names are:

```
# ipcs -l

----- Shared Memory Limits -----
max number of segments = 4096           // SHMMNI
max seg size (kbytes) = 32768           // SHMMAX
max total shared memory (kbytes) = 8388608 // SHMALL
min seg size (bytes) = 1

----- Semaphore Limits -----
max number of arrays = 1024             // SEMMNI
max semaphores per array = 250          // SEMMSL
max semaphores system wide = 256000    // SEMMNS
max ops per semop call = 32             // SEMOPM
semaphore max value = 32767

----- Messages: Limits -----
max queues system wide = 1024           // MSGMNI
max size of message (bytes) = 65536     // MSGMAX
default max size of queue (bytes) = 65536 // MSGMNB
```

- Beginning with the first section on Shared Memory Limits, the **SHMMAX** limit is the maximum size of a shared memory segment on a Linux system. The **SHMALL** limit is the maximum allocation of shared memory pages on a system.
 - It is recommended to set the **SHMMAX** value to be equal to the amount of physical memory on your system. However, the minimum required on x86 systems is 268435456 (256 MB) and for 64-bit systems, it is 1073741824 (1 GB).
 - The **SHMALL** parameter is set to 8 GB by default (8388608 KB = 8 GB). If you have more physical memory than 8 GB, and it is to be used for DB2, then this parameter increases to approximately 90% of your computer's physical memory. For instance, if you have a computer system with 16 GB of memory to be used primarily for DB2, then **SHMALL** should be set to 3774873 (90% of 16 GB is 14.4 GB; 14.4 GB is then divided by 4 KB, which is the base page size). The **ipcs** output converted **SHMALL** into kilobytes. The kernel requires this value as a number of pages. If you are upgrading to DB2 Version 10.1 and you are not using the default **SHMALL** setting, you must increase the **SHMALL** setting by an additional 4 GB. This increase in memory is required by the fast communication manager (FCM) for additional buffers or channels.
- The next section covers the amount of semaphores available to the operating system. The kernel parameter **sem** consists of four tokens, **SEMMSL**, **SEMMNS**, **SEMOPM** and **SEMMNI**. **SEMMNS** is the result of **SEMMSL** multiplied by **SEMMNI**. The database manager requires that the number of arrays (**SEMMNI**) be increased as necessary. Typically, **SEMMNI** should be twice the maximum number of agents expected on the system multiplied by the number of logical partitions on the database server computer plus the number of local application connections on the database server computer.

- The third section covers messages on the system.
 - The **MSGMNI** parameter affects the number of agents that can be started; the **MSGMAX** parameter affects the size of the message that can be sent in a queue, and the **MSGMNB** parameter affects the size of the queue.
 - The **MSGMAX** parameter should be changed to 64 KB (that is, 65536 bytes), and the **MSGMNB** parameter should be increased to 65536.
3. Modify the kernel parameters that you have to adjust by editing the `/etc/sysctl.conf` file. If this file does not exist, create it. The following lines are examples of what should be placed into the file:


```
#Example for a computer with 16GB of RAM:
kernel.shmmni=4096
kernel.shmmax=17179869184
kernel.shmall=8388608
#kernel.sem=<SEMMSL> <SEMMNS> <SEMOPM> <SEMMNI>
kernel.sem=250 1024000 32 4096
kernel.msgmni=16384
kernel.msgmax=65536
kernel.msgmnb=65536
```
 4. Run **sysctl** with **-p** parameter to load in **sysctl** settings from the default file `/etc/sysctl.conf`:


```
sysctl -p
```
 5. Optional: Have the changes persist after every reboot:
 - (SUSE Linux) Make `boot.sysctl` active.
 - (Red Hat) The `rc.sysinit` initialization script reads the `/etc/sysctl.conf` file automatically.

For the latest information about supported Linux distributions, see <http://www.ibm.com/software/data/db2/linux/validate/>.

Kernel parameter requirements (Linux)

The database manager uses a formula to automatically adjust kernel parameter settings and eliminate the need for manual updates to these settings.

Interprocess communication kernel parameters

When instances are started, if an interprocess communication (IPC) kernel parameter is below the enforced minimum value, the database manager updates it to enforced minimum value. The IPC kernel parameter values changed when a DB2 instance is started do not persist when the system is rebooted. After a reboot, kernel settings might be lower than the enforced values until a DB2 instance is started. By adjusting any kernel parameter settings, the database manager prevents unnecessary resource errors.

For the latest information about supported Linux distributions, see <http://www.ibm.com/software/data/db2/linux/validate/>.

Table 10. Enforced minimum settings for Linux interprocess communication kernel parameters

IPC kernel parameter	Enforced minimum setting
kernel.shmmni (SHMMNI)	256 * <size of RAM in GB>
kernel.shmmax (SHMMAX)	<size of RAM in bytes> ¹
kernel.shmall (SHMALL)	2 * <size of RAM in the default system page size> ²
kernel.sem (SEMMNI)	256 * <size of RAM in GB>

Table 10. Enforced minimum settings for Linux interprocess communication kernel parameters (continued)

IPC kernel parameter	Enforced minimum setting
<code>kernel.sem (SEMMSL)</code>	250
<code>kernel.sem (SEMMNS)</code>	256 000
<code>kernel.sem (SEMOPM)</code>	32
<code>kernel.msgmni (MSGMNI)</code>	1 024 * <size of RAM in GB>
<code>kernel.msgmax (MSGMAX)</code>	65 536
<code>kernel.msgmnb (MSGMNB)</code>	65 536 ³
<ol style="list-style-type: none"> 1. On 32-bit Linux operating systems, the enforced minimum setting for SHMMAX is limited to 4 294 967 295 bytes. 2. SHMALL limits the total amount of virtual shared memory that can be allocated on a system. Each DB2 data server efficiently manages the amount of system memory it consumes, also known as committed memory. The DB2 data server allocates more virtual memory than it commits to support memory preallocation and dynamic memory management. Memory preallocation benefits performance. Dynamic memory management is the process of growing and shrinking real memory usage within separate virtual shared memory areas. To support memory preallocation and dynamic memory management effectively data servers frequently have to allocate more virtual shared memory on a system than the amount of physical RAM. The kernel requires this value as a number of pages. 3. Load performance might benefit from a larger message queue size limit, specified in bytes by MSGMNB. You can view message queue usage can by running the ipcs -q command. If the message queues are at capacity, or reaching capacity, during load operations, consider increasing the number of bytes the message queue size limit. 	

Other recommended kernel parameter settings

Other recommended kernel parameter settings are listed in the following table.

Table 11. Configuring other Linux kernel parameters

Recommended kernel parameter setting	Configuring the kernel parameters for DB2 data server
<code>vm.swappiness=0</code>	This parameter defines how prone the kernel is to swapping application memory out of physical random access memory (RAM). The default setting is <code>vm.swappiness=60</code> . The recommended kernel parameter setting, <code>vm.swappiness=0</code> , configures the kernel to give preference to keeping application memory in RAM instead of assigning more memory for file caching. This setting avoids unnecessary paging and excessive use of swap space. This setting is especially important for data servers configured to use the self-tuning memory manager (STMM).

Table 11. Configuring other Linux kernel parameters (continued)

Recommended kernel parameter setting	Configuring the kernel parameters for DB2 data server
vm.overcommit_memory=0	This parameter influences how much virtual memory the kernel permits allocating. The default setting, vm.overcommit_memory=0 , sets the kernel to disallow individual processes from making excessively large allocations, however the total allocated virtual memory is unlimited. Having unlimited virtual memory is important for DB2 data servers, which retain additional unused virtual memory allocations for dynamic memory management. Unreferenced allocated memory is not backed by RAM or paging space on Linux systems. Avoid setting vm.overcommit_memory=2 , as this setting limits the total amount of virtual memory that can be allocated, which can result in unexpected errors.

Installation requirements for DB2 servers and IBM data server clients (Solaris)

Before you install a DB2 database product on Solaris, ensure that your system meets the operating system, hardware, and software requirements. The **db2prereqcheck** command checks whether your system meets the prerequisites for the installation.

DB2 database products are available in several different editions. Some DB2 database products and features are only available on specific operating systems.

Table 12. Solaris installation requirements

Operating System	Hardware
Solaris 10 <ul style="list-style-type: none"> • 64-bit kernel • Update 9 	UltraSPARC or SPARC64 processors
Solaris 10 <ul style="list-style-type: none"> • 64-bit kernel • Update 9 	Solaris x64 (Intel 64 or AMD64)

- Support is only for the DB2 product to be installed on either local zone or global zone. When installed in the global zone, it can be shared by other zones.

Kernel configuration considerations

The kernel configuration parameters are set in `/etc/system`. If the kernel parameter being modified is not listed as dynamic, a system reboot is required to make the changes to `/etc/system` take effect. These parameters must be set before you install an IBM data server client.

Software considerations

- `libpam.so.0` (32-bit) is required for DB2 database servers to run 32-bit non-SQL routines.

- (Client and server) To use Kerberos authentication, install the krb5 Kerberos support via the available Kerberos packages.
- One of the following browsers is required to view online help and to run First Steps (**db2fs**):
 - Firefox 3.0 and later
 - Google Chrome
 - Safari 4.0
- An X Window System software capable of rendering a graphical user interface is required if you want to use the DB2 Setup wizard to install a DB2 database product
- For details regarding known Solaris issues, see www.ibm.com/support/docview.wss?uid=swg21257606

Security patches can be obtained from the <http://java.sun.com> website. From the SunSolve online website, click **Patches** in the left panel.

The Java2 Standard Edition (J2SE) Solaris Operating System Patch Clusters and the SUNWlibC software are also required and can be obtained from the <http://java.sun.com> website.

The Fujitsu PRIMEPOWER patches for the Solaris Operating System can be downloaded from FTSI at: <http://download.ftsi.fujitsu.com/>.

DB2 database products support the following Solaris concepts:

- Solaris Logical Domains (LDoms)
- Solaris Zones
- ZFS filesystems

For details about virtualization technology supported by DB2 products, see <http://www.ibm.com/developerworks/wikis/display/im/DB2+Virtualization+Support>.

Mounting considerations

For root installations of DB2 products on UNIX operating systems, do not mount your file system with the *noexec* option.

Installation requirements for DB2 servers and IBM data server clients (Windows)

Before you install a DB2 database product on Windows operating systems, ensure that the system you choose meets the necessary minimum operating system, hardware, and software requirements.

Table 13. Windows workstation platforms

Operating System	Prerequisites	Hardware
Windows XP Professional (32-bit and 64-bit)	Windows XP Service Pack 2 or later	All Intel and AMD processors capable of running the supported Windows operating systems (32-bit and 64-bit based systems)
Windows Vista Business (32-bit and 64-bit)	IBM Data Server Provider for .NET client applications and CLR server-side procedures require .NET 2.0 or later framework runtime.	
Windows Vista Enterprise (32-bit and 64-bit)		
Windows Vista Ultimate (32-bit and 64-bit)		
Windows 7 Professional (32-bit and 64-bit)	64-bit IBM data server provider for .NET applications are supported.	
Windows 7 Enterprise (32-bit and 64-bit)	All Windows Vista and Windows 7 service packs are supported.	
Windows 7 Ultimate (32-bit and 64-bit)		
Windows 8 Standard (32-bit and 64-bit) and Windows 8 Professional (32-bit and 64-bit)		

Table 14. Windows server platforms

Operating System	Prerequisites	Hardware
Windows 2003 Datacenter Edition (32-bit and 64-bit)	Service Pack 2 or later.	All Intel and AMD processors capable of running the supported Windows operating systems (32-bit and 64-bit based systems).
Windows 2003 Enterprise Edition (32-bit and 64-bit)	R2 is also supported	
Windows 2003 Standard Edition (32-bit and 64-bit)	IBM data server provider for .NET client applications and CLR server-side procedures require .NET 2.0 or later framework runtime. 64-bit IBM data server provider for .NET applications are supported.	
Windows Server 2008 Datacenter Edition (32-bit and 64-bit) and Windows Server 2008 R2 (64-bit)	IBM data server provider for .NET client applications and CLR server-side procedures require .NET 2.0 or later framework runtime.	
Windows Server 2008 Enterprise Edition (32-bit and 64-bit) and Windows Server 2008 R2 (64-bit)	64-bit IBM data server provider for .NET applications are supported.	
Windows Server 2008 Standard Edition (32-bit and 64-bit) and Windows Server 2008 R2 (64-bit)	All Windows Server 2008 service packs are supported.	
Windows Server 2012 Datacenter Edition (32-bit and 64-bit) and Windows Server 2012 Standard Edition (32-bit and 64-bit)		
DB2 V9.7.0.2 is Windows Server 2008 R2 certified. See http://www.windowsservercatalog.com/results.aspx?text=db2&bCatID=1282&avc=10&ava=0&OR=5&Go for a complete list of Windows-certified DB2 products.		

Note: DB2 database products support the hardware-enforced Data Execution Prevention (DEP) feature that is built into some Windows operating systems.

Additional software considerations

- Windows Installer 3.0 is required. It will be installed by the installer if it is not detected.
- IBM Data Server Provider for .NET client applications and CLR server-side procedures require .NET 2.0 or later framework runtime. In an x64 environment, 32-bit IBM data server provider for .NET applications will run in the WOW64 emulation mode.
- If you plan to use LDAP (Lightweight Directory Access Protocol), you should use either a Microsoft LDAP client or the IBM Tivoli Directory Server v6 client (also known as the IBM LDAP client which is included with DB2 database products). Before installation of the Microsoft Active Directory, you will need to extend your directory schema using the **db2schex** utility, which can be found on the installation media under the db2\Windows\utilities directory.

The Microsoft LDAP client is included with Windows operating systems.

- One of the following browsers is required to view online help, run the DB2 install launchpad (setup.exe), and to run First Steps (**db2fs**):
 - Firefox 2.0 and later
 - Internet Explorer 6.0 and later
 - Mozilla-based browsers 1.7 and later
 - Safari 3.0 and later

Setting up Network Time Protocol

Ensure that all servers in a DB2 instance have similar time settings by setting up a Network Time Protocol (NTP) server and clients.

Configuring hosts as Network Time Protocol clients

This topic outlines how to configure a Network Time Protocol (NTP) server to ensure that all servers that the instance runs on have similar time settings.

About this task

By default, the NTP server is setup and configuration is handled by the DB2 installer.

Procedure

To configure hosts as NTP clients:

1. Specify the NTP server to synchronize to. Open the `/etc/ntp.conf` file and add the following lines:

```
server ntp_server_hostname
driftfile /etc/ntp.drift
tracefile /etc/ntp.trace
```

where `ntp_server_hostname` represents the hostname or IP address of an NTP server. If you do not have access to an existing NTP server, refer to the “Setting up an NTP Server” topic.

2. Configure the NTP server to start at system restart:
 - On AIX operating systems, edit the `/etc/rc.tcpip` file and uncomment the following line:

```
start /usr/sbin/xntpd -x "$src_running"
```
 - On Linux operating systems, run the **chkconfig ntp on** command.
3. Start the NTP server:
 - On AIX operating systems, run the **startsrc -s xntpd** command.
 - On Linux operating systems, run the **server ntp start** command.
4. Verify that the daemon is synchronized. If the daemon is significantly not synchronized with the system clock, it can take over 10 minutes for it to become synchronized. To ensure the daemon is synchronized before continuing you can check the “sys stratum” field:
 - On AIX operating systems, use the output of the `lssrc -ls xntpd` command.
 - On Linux operating systems, use the output of the `ntptrace` command.

The `sys stratum` field should be less than 16 to continue to the next step.

5. Finish the synchronization process by running the `ntpdate -d ntp_server_hostname` command.

Setting up a Network Time Protocol server

This topic outlines how to setup a Network Time Protocol (NTP) server to ensure that all servers that the instance runs on have similar time settings.

By default, the NTP server is setup by the DB2 installer. During instance creation and update, the NTP is automatically setup on the primary CF. This setup is a best practice that will help with problem diagnosis and informal server monitoring.

Before you begin

In a DB2 pureScale environment, to synchronize operations and facilitate time sensitive operations, the system clocks of all members must be synchronized.

For this reason, NTP must be setup on all of your members. The DB2 installer installs and sets up the NTP server and clients. By default, the NTP server is the primary CF. You can use the steps in this topic to change that default setting and assign another host to act as the NTP server.

The DB2 database manager checks whether NTP is correctly set up on each member and verifies the presence of the NTP daemon. The DB2 database manager also periodically compares time stamps between members to check that they differ by less than the value set by the MAX_TIME_DIFF database manager configuration parameter.

In a DB2 pureScale environment, the value of MAX_TIME_DIFF represents the maximum number of minutes allowed between the system clocks of two members; the default is 1 minute. For each member, if NTP is not correctly set up or the NTP daemon is not present, or if time stamps between members differ by more than the value of MAX_TIME_DIFF, warnings are written to the **db2diag** log files.

About this task

If you do not have access to an NTP time server, use the following steps to configure one of the hosts in the cluster or another host accessible by members as an NTP server. The system clock on this host will be used as the clock that is synchronized. This host must remain online to continue synchronizing the clocks on the hosts.

Procedure

1. Choose the host that you will setup as an NTP server.
2. Setup the NTP daemon to synchronize to its own system clock. Edit `/etc/ntp.conf` and add the following lines:

```
server 127.127.1.0 prefer # the logical IP address for the NTP server to sync
to its own system clock
fudge 127.127.1.0
driftfile /etc/ntp.drift
tracefile /etc/ntp.trace
```
3. Configure the NTP server to start at system restart.
 - On AIX operating systems, edit the file `/etc/rc.tcpip` and uncomment the following entry:

```
start /usr/sbin/xntpd -x "$src_running"
```
 - On Linux operating systems, run the **chkconfig ntp on** command.
4. Start the NTP server:
 - On AIX operating systems, run the `startsrc -s xntpd` command.

- On Linux operating systems, run the **service ntp start** command.

Directory structure for your installed DB2 database product (Windows)

When you install DB2 database products, you can specify a DB2 database product installation path or else use the default path. After installation, DB2 objects are created in these directories.

Follow these steps to verify the DB2 product you have installed on Windows.

1. From a command prompt, type the **regedit** command. The Registry Editor window opens.
2. Expand **HKEY_LOCAL_MACHINE > Software > IBM > DB2**

The DB2 product you have installed will be displayed.

The following table shows the location of DB2 objects after a default installation.

Table 15. DB2 objects and their locations

DB2 Object	Location
DAS information	<ul style="list-style-type: none"> • For Windows XP and Windows 2003 operating systems: C:\Documents and Settings\All Users\Application Data\IBM\DB2\DB2COPY1 • For Windows Vista and later operating systems: C:\ProgramData\IBM\DB2\DB2COPY1
Database configuration file SQLDBC0N	C:\DB2\NODE0000\SQL00001
Database directory Contains files needed for: <ul style="list-style-type: none"> • buffer pool information • history information • log control files • storage path information • table space information 	C:\DB2\NODE0000\SQL00001
Database manager configuration file db2system	<ul style="list-style-type: none"> • For Windows XP and Windows 2003 operating systems: C:\Documents and Settings\All Users\Application Data\IBM\DB2\DB2COPY1\DB2 • For Windows Vista and later operating systems: C:\ProgramData\IBM\DB2\DB2COPY1\DB2
DB2 commands	C:\Program Files\IBM\SQLLIB\BIN
DB2 error messages file db2diag log files	<ul style="list-style-type: none"> • For Windows XP and Windows 2003 operating systems: C:\Documents and Settings\All Users\Application Data\IBM\DB2\DB2COPY1\DB2 • For Windows Vista and later operating systems: C:\ProgramData\IBM\DB2\DB2COPY1\DB2
DB2 installation path	C:\Program Files\IBM\SQLLIB
Directory for event monitor data	C:\DB2\NODE0000\SQL00001\DB2EVENT

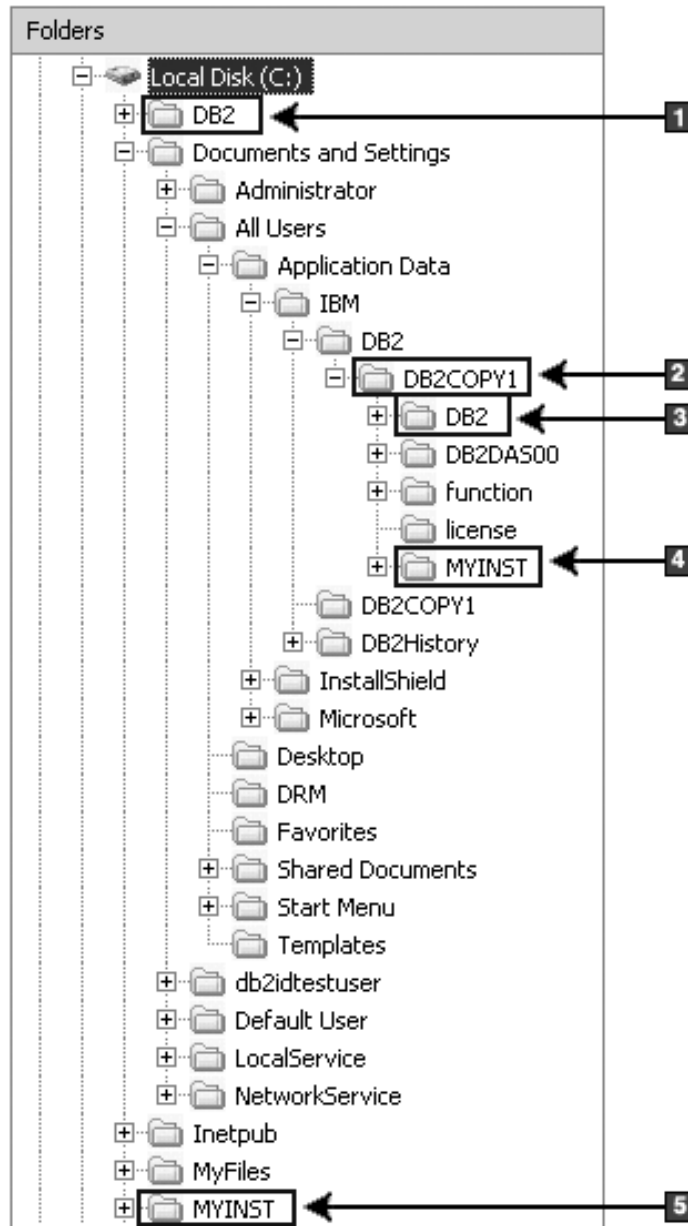
Table 15. DB2 objects and their locations (continued)

DB2 Object	Location
Directory for transaction log files	C:\DB2\NODE0000\SQL00001\LOGSTREAM0000
Installation log file	<ul style="list-style-type: none"> For Windows XP and Windows 2003 operating systems: C:\Documents and Settings\Administrator\My Documents\DB2LOG For Windows Vista and later operating systems: C:\Users\USER_NAME\Documents\DB2LOG
Instance	<ul style="list-style-type: none"> For Windows XP and Windows 2003 operating systems: C:\Documents and Settings\All Users\Application Data\IBM\DB2\DB2COPY1\DB2 For Windows Vista and later operating systems: C:\ProgramData\IBM\DB2\DB2COPY1\DB2
Instance information	<ul style="list-style-type: none"> For Windows XP and Windows 2003 operating systems: C:\Documents and Settings\All Users\Application Data\IBM\DB2\DB2COPY1\DB2 For Windows Vista and later operating systems: C:\ProgramData\IBM\DB2\DB2COPY1\DB2
Node directory	<ul style="list-style-type: none"> For Windows XP and Windows 2003 operating systems: C:\Documents and Settings\All Users\Application Data\IBM\DB2\DB2COPY1\DB2\SQLNODIR For Windows Vista and later operating systems: C:\ProgramData\IBM\DB2\DB2COPY1\DB2\SQLNODIR
Local database directory for the instance called DB2	C:\DB2\NODE0000\SQLDBDIR
Partitioned database environment file db2nodes.cfg	<ul style="list-style-type: none"> For Windows XP and Windows 2003 operating systems: C:\Documents and Settings\All Users\Application Data\IBM\DB2\DB2COPY1\DB2 For Windows Vista and later operating systems: C:\ProgramData\IBM\DB2\DB2COPY1\DB2
System database directory	<ul style="list-style-type: none"> For Windows XP and Windows 2003 operating systems: C:\Documents and Settings\All Users\Application Data\IBM\DB2\DB2COPY1\DB2\SQLDBDIR For Windows Vista and later operating systems: C:\ProgramData\IBM\DB2\DB2COPY1\DB2\SQLDBDIR

The following figures illustrate an example of the DB2 directory structure after installation on Windows XP or Windows 2003 operating systems using the default

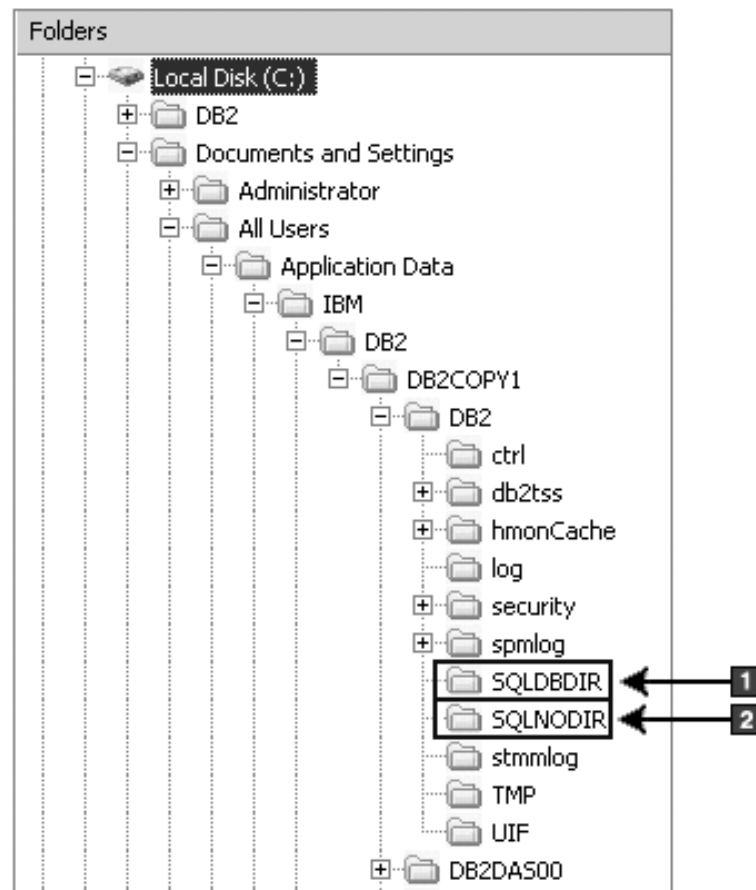
options. In these figures, there are two instances, DB2 and MYINST. The directories DB2 and MYINST under the local disk C: will only appear if a database has been created under the appropriate instance.


Directory structure - instance information



1. Contains the databases created under the C: drive for the instance named DB2.
2. Contains the information for the DAS.
3. Contains the instance information for the instance named DB2.
4. Contains the instance information for the instance named MYINST.
5. Contains the databases created under the C: drive for the instance named MYINST.


Directory structure - directory information






1. System database directory
2. Node directory
3. The db2diag log files DB2 error messages.
4. The db2nodes.cfg file is used in a partitioned database environment.
5. Database manager configuration file

Directory structure - local directory information



1. Local database directory for the instance DB2
2. Local database directory for the instance MYINST
3. Database configuration file

Directory structure - installation log file location



Directory structure for your installed DB2 database product (Linux)

During a root installation, you can specify where the subdirectories and files for the DB2 database product will be created. For non-root installations, you cannot choose where DB2 products are installed; you must use the default locations.

Note: For non-root installations, all DB2 files (program files and instance files) are located in or beneath the `$HOME/sqllib` directory, where `$HOME` represents the non-root user's home directory.

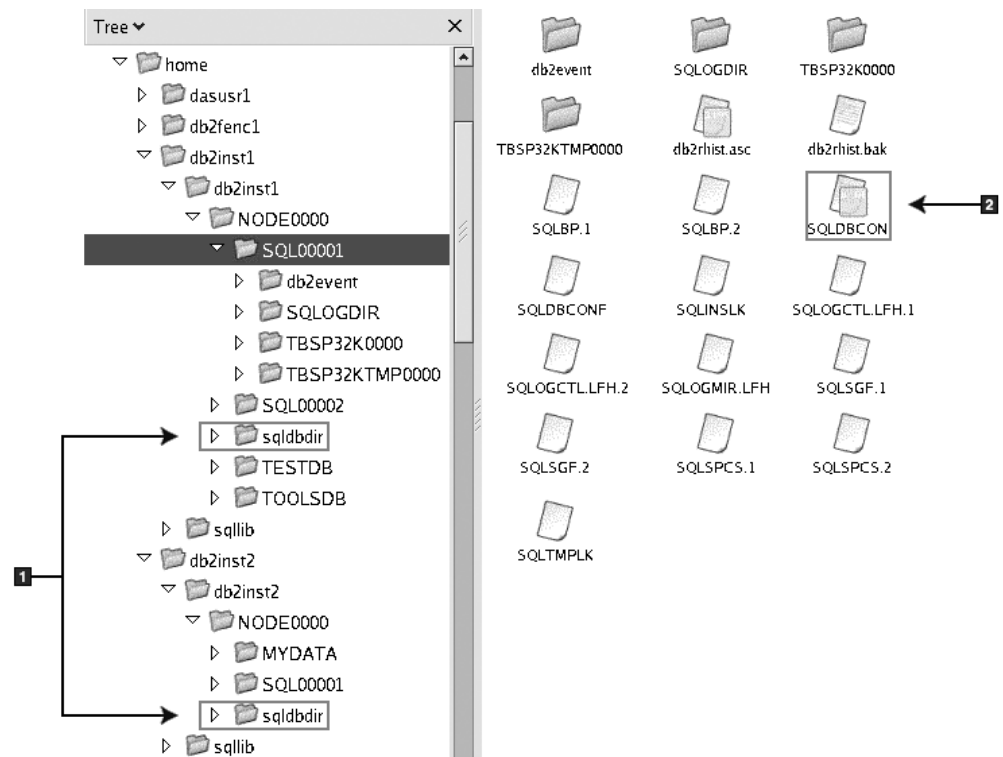
After installation, the DB2 objects are created in various directories. The following table shows the location of DB2 objects after a default root installation.

Table 16. Location of DB2 objects after a default root installation

DB2 Object	Location
DAS home directory	home/dasusr1
DAS information	home/dasusr1/das
Database configuration file SQLDBCON	home/db2inst1/db2inst1/NODE0000/SQL00001
Database directory Contains files needed for: <ul style="list-style-type: none"> • buffer pool information • history information • log control files • storage path information • table space information 	home/db2inst1/db2inst1/NODE0000/SQL00001
Database manager configuration file db2system	home/db2inst1/sqllib
DB2 commands	/opt/IBM/db2/V10.1/bin
DB2 error messages file (db2diag log file)	home/db2inst1/sqllib/db2dump
DB2 installation path	default is /opt/IBM/db2/V10.1
Directory for event monitor data	home/db2inst1/db2inst1/NODE0000/SQL00001/db2event
Directory for transaction log files	home/db2inst1/db2inst1/NODE0000/SQL00001/LOGSTREAM0000
Installation log file db2install.history	/opt/IBM/db2/V10.1/install/logs
Instance home directory	home/db2inst1
Instance information	home/db2inst1/sqllib
Local database directory for the instance	home/db2inst1/db2inst1/NODE0000/sqldbdir
Partitioned database environment file db2nodes.cfg	home/db2inst1/sqllib
System database directory	home/db2inst1/sqllib/sqldbdir

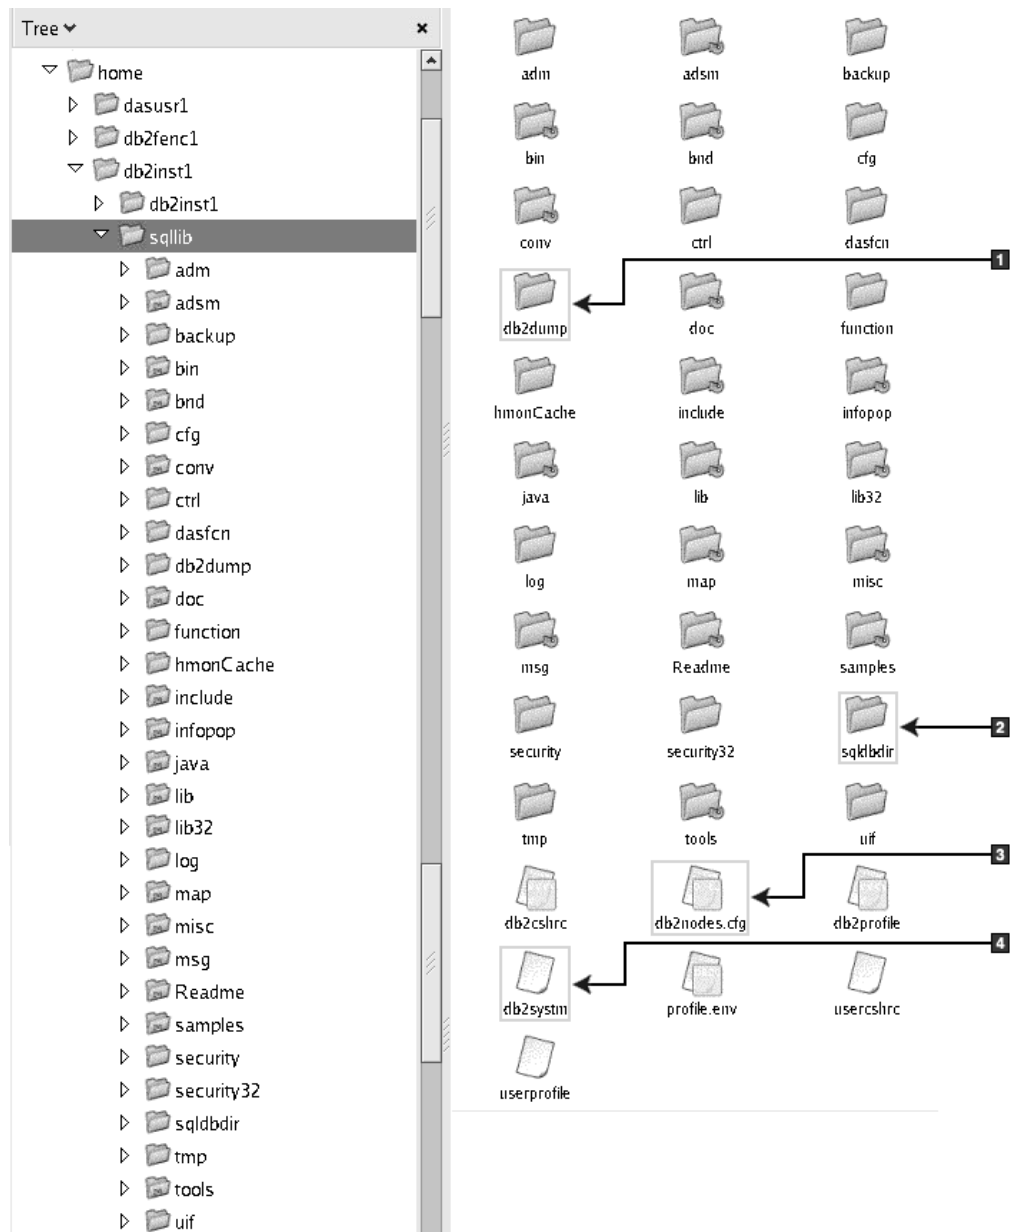
The following figures illustrate an example of the DB2 directory structure after a root installation. In these examples, there are two instances, db2inst1 and db2inst2.

Directory structure - default local database directory information for the DB2 instance db2inst1



1. Local database directories.
2. Database configuration file


Directory structure - directory information for the DB2 instance db2inst1



1. The db2diag log files in this directory logs DB2 error messages.
2. System database directory
3. The db2nodes.cfg file is used in a partitioned database environment.
4. Database manager configuration file

The dasusr1 directory contains the DB2 administration server (DAS) configuration files and the symbolic links to DB2 core files.

Directory structure - install log file location



This figure illustrates the location of your install log file. If there were multiple installations under the same installation path, the `db2install.history` file will be indexed as `db2install.history.n` where *n* represents a four digit number, for example, 0000, or 0001.

The DB2 installation directory `/opt/IBM/db2/V10.1` contains the installed DB2 files.

DB2 Setup wizard language support

Before running the DB2 Setup wizard, review the language support considerations for the wizard.

Supported DB2 interface languages

DB2 language support for DB2 interfaces can be categorized into server group languages and client group languages.

Server group languages will translate most messages, help, and DB2 graphical interface elements. Client group languages will translate the IBM Data Server Runtime Client component, which will include most messages and certain help documentation.

Server group languages include: Brazilian Portuguese, Czech, Danish, Finnish, French, German, Italian, Japanese, Korean, Norwegian, Polish, Russian, Simplified Chinese, Spanish, Swedish, and Traditional Chinese.

Client group languages include: Arabic, Bulgarian, Croatian, Dutch, Greek, Hebrew, Hungarian, Portuguese, Romanian, Slovak, Slovenian, and Turkish.

Do not confuse languages supported by the DB2 database product with languages supported by the DB2 interface. Languages supported by the DB2 database product means the languages in which data can exist. These languages are a superset of languages supported by the DB2 interface.

Displaying the DB2 Setup wizard in your national language (Linux and UNIX)

The **db2setup** command queries the operating system to determine the existing language settings. If the language setting of your operating system is supported by **db2setup**, then that language will be used when displaying the DB2 Setup wizard.

If your system uses the same code pages but different locale names than those supported by the DB2 interface, you can still see the translated **db2setup** by setting your **LANG** environment variable to the appropriate value by entering the following command:

bourne (sh), korn (ksh), and bash shells:

```
LANG=locale
export LANG
```

C shell:

```
setenv LANG locale
```

where *locale* is a locale supported by the DB2 interface.

Language identifiers for running the DB2 Setup wizard in another language

If you want to run the DB2 Setup wizard in a language different from the default language on your computer, you can start the DB2 Setup wizard manually, specifying a language identifier. The language must be available on the platform where you are running the installation.

On Windows operating systems, you can run **setup.exe** with the **-i** parameter to specify the two-letter language code of the language the installation is to use.

On Linux and UNIX operating systems, it is recommended that you set the **LANG** environment variable to display the DB2 Setup wizard in your national language.

Table 17. Language identifiers

Language	Language identifier
Arabic (available on Windows platforms only)	ar

Table 17. Language identifiers (continued)

Language	Language identifier
Brazilian Portuguese	br
Bulgarian	bg
Chinese, Simplified	cn
Chinese, Traditional	tw
Croatian	hr
Czech	cz
Danish	dk
Dutch	nl
English	en
Finnish	fi
French	fr
German	de
Greek	el
Hungarian	hu
Indonesian (available on Windows platforms only)	id
Italian	it
Japanese	jp
Korean	kr
Lithuanian (available on Windows platforms only)	lt
Norwegian	no
Polish	pl
Portuguese	pt
Romanian	ro
Russian	ru
Slovak	sk
Slovenian	sl
Spanish	es
Swedish	se
Turkish	tr

Changing the DB2 interface language (Linux and UNIX)

The interface language of the DB2 database product is the language that appears in messages, help, and graphical tool interfaces.

Before you begin

Do not confuse languages supported by the DB2 database product with languages supported by the DB2 interface. Languages supported by the DB2 database product, that is, languages that *data* can exist in, are a superset of languages supported by the DB2 interface.

Support for the DB2 interface language you want to use must be installed on your system. DB2 interface language support is selected and installed when you install a DB2 database product using the DB2 Setup wizard. If you change the interface language of the DB2 database product to a supported interface language that has not been installed, the DB2 interface language will default to the operating system language. If the operating system language is not supported, English is used as the DB2 interface language.

DB2 interface language support is selected and installed when you install your DB2 database product using the DB2 Setup wizard or by using the National Language Package.

About this task

To check which public locales are available in your system, run the **\$ locale -a** command.

Procedure

To change the DB2 interface language:

Set the **LANG** environment variable to the locale you want.

- For bourne (sh), korn (ksh), and bash shells:

```
LANG=locale
export LANG
```

- For C shell:

```
setenv LANG locale
```

For example, to interface with the DB2 database product in French, you must have the French language support installed and you must set the **LANG** environment variable to a French locale, for example, `fr_FR`.

Changing the DB2 product interface language (Windows)

The DB2 interface language is the language that appears in messages, help, and graphical tool interfaces.

About this task

Do not confuse languages supported by a DB2 database product with languages supported by the DB2 interface. Languages supported by a DB2 database product means the languages in which data can exist. These languages are a superset of languages supported by the DB2 interface.

The DB2 interface language you want to use must be installed on your system. The DB2 database product interface languages are selected and installed when you install a DB2 database product using the DB2 Setup wizard. If you change the interface language of a DB2 database product to a supported interface language that has not been installed, the DB2 database product interface language will default to the operating system language first, and if that is not supported, English.

Changing the interface language for a DB2 database product on Windows requires that you change the default language setting for your Windows operating system.

Procedure

To change the DB2 database product interface language on Windows operating systems:

1. Through the Control Panel, select **Regional and Language Options**.
2. On the **Regional Options** tab under **Standards and formats**, select the appropriate language. On Windows 2008 and Windows Vista or higher, use the **Formats** tab for this step.
3. On the **Regional Options** tab under **Location**, select the location that corresponds to the appropriate language.
4. On the **Advanced** tab under **Language for non-Unicode programs** select the appropriate language. On Windows 2008 and Windows Vista or higher, on the **Administrative** tab, under **Language for non-unicode programs**, click **Change system locale** and select the appropriate language. You will then be asked to reboot, click **Cancel**.
5. On the **Advanced** tab under **Default user account settings**, check the **Apply all settings to the current user account and to the default user profile** box. On Windows 2008 and Windows Vista or higher, on the **Administrative** tab under **reserved accounts**, click **Copy to reserved accounts** and check the accounts that you want to copy the language settings to.
6. You will be asked to reboot before these changes come into effect.

What to do next

Refer to your operating system help for additional information about changing the default system language.

Password rules

When choosing a password, you need to be aware that your password length can be up to the maximum number of characters supported by your operating system.

Some operating systems might have more password rules, such as minimum length and simplicity, and some operating systems can be configured to use different password encryption algorithms. See the appropriate operating system documentation for more information.

Note: To enable long password support on the AIX operating system, Version 6.1, or later, install APAR IZ35001 on the DB2 client and server machines.

Multiple DB2 copies on the same computer (Windows)

You can use multiple DB2 copies on the same computer. Each DB2 copy can be at the same or different code levels.

The benefits of this configuration include:

- The ability to run applications that require different DB2 versions on the same host at the same time.
- The ability to run independent copies of DB2 products for different functions.
- The ability to test on the same computer before moving the production database to a later version of the DB2 product.
- For independent software vendors, the ability to embed a DB2 server product into your product and hide the DB2 database from your users.

A DB2 copy is a group of DB2 products that are installed at the same location.

Differences when only one DB2 copy is installed

- During installation, a unique default DB2 copy name is generated. You can change the name of the default DB2 copy as you go through the DB2 Setup wizard. You cannot change the DB2 copy name after the installation is completed.
- Applications use the default DB2 copy in an environment similar to the DB2 Version 8 environment.

Differences when multiple DB2 copies are installed on the same computer

- DB2 Version 10.1 copies can coexist with copies of DB2 Version 9.7 or earlier releases with the following restrictions.
 - You must define a default copy. After DB2 Version 10.1 is installed, you can use the Global Switcher to set the Version 10.1 copy as the default copy.
 - Optional: You can configure each DB2 copy to use a different DB2 Information Center.


Note: You can have only one copy of the DB2 Information Center installed on the same system at the same Release level. However, you cannot have one DB2 Information Center at a different fix pack level on the same host. You can however configure the DB2 database server to access these DB2 Information Centers remotely.

- Only the IBM Data Server Provider for .NET from the default IBM database client interface copy is registered in the Global Assembly Cache.
 - Each DB2 copy must have unique instance names. For a silent installation with the **NO_CONFIG** flag set to YES, the default instance is not created. However, when you create the instance after the installation, it must be unique. The default name of the instance is “DB2”. If an instance of the “DB2” name exists, a unique name for the instance is generated. The unique name is generated by using the “DB2” name and adding an underscore and generating the last two characters in numeric sequence. The subsequent instance names generated are “DB2_01”, “DB2_02”, and so on.
- For Microsoft COM+ applications, use and distribute the IBM Data Server Driver Package (installer) or IBM Data Server Driver for ODBC and CLI (compressed file) with your application instead of the IBM Data Server Runtime Client. Only one Data Server Runtime Client can be used for COM+ applications at a time. The IBM Data Server Driver Package (installer) or IBM Data Server Driver for ODBC and CLI (compressed file) does not have this restriction. Microsoft COM+ applications accessing DB2 data sources are only supported with the default DB2 copy. Concurrent support of COM+ applications accessing different DB2 copies is not supported. You can change the default DB2 copy with the Default DB2 Copy Selection Wizard.


Choosing a default when installing a new DB2 copy

The following figure illustrates a scenario where you have installed multiple DB2 copies: *DB2COPY1*, *DB2COPY2*, and on to *DB2COPYn*. One of the DB2 copies is selected by you as the default DB2 copy. In this case, *DB2COPY1* is selected as the default DB2 copy.

System environment

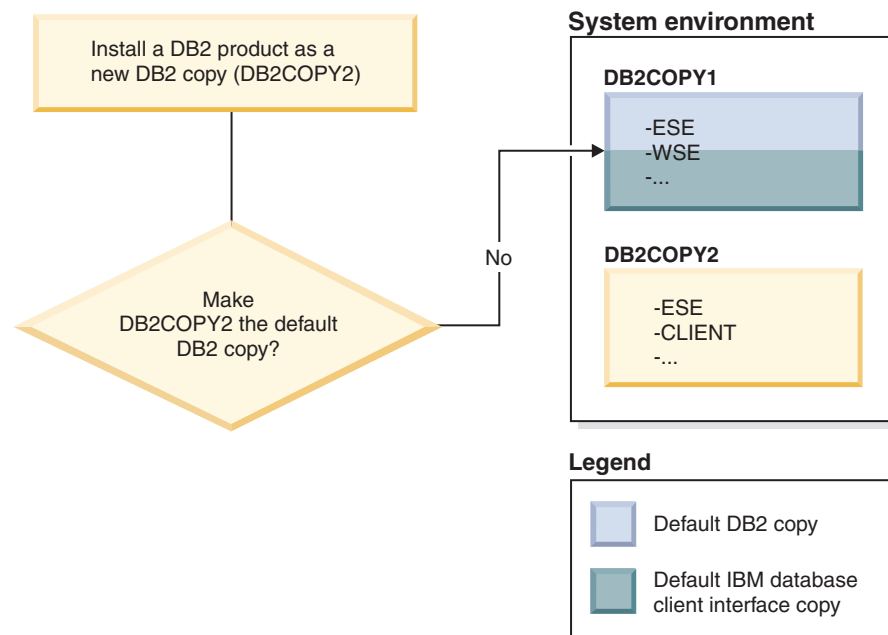


Legend



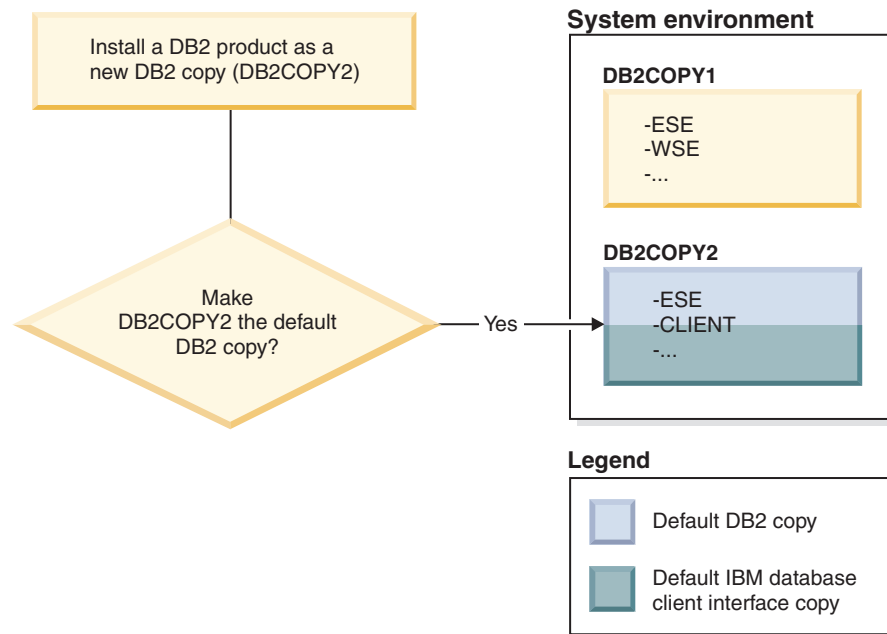
Now imagine a scenario where you install one DB2 copy and you call it *DB2COPY1*. *DB2COPY1* is the default DB2 copy and the default IBM database client interface copy.

Then you install a DB2 product in a new DB2 copy called *DB2COPY2*. During the installation of *DB2COPY2*, you are asked if you want to make the new DB2 copy the default DB2 copy. If you respond “No”, then *DB2COPY1* remains the default DB2 copy and the default IBM database client interface copy as shown in the following figure:



However, consider the same scenario but you respond “Yes” when asked if you want to make the new DB2 copy the default DB2 copy. In this case, *DB2COPY2* becomes the new default DB2 copy and the default IBM database client interface

copy as shown in the following figure:



Applications that run as a service

Applications that dynamically bind DB2 DLL files, for example applications that are linked with `db2api.lib`, find the DB2 DLL files in the `PATH`. This means that existing applications that were not developed for multiple DB2 versions use the default DB2 copy. To work around this behavior, the application can use the `db2SelectDB2Copy` API before loading any DB2 libraries.

Note: When linking with `db2api.lib`, the functions resolve to different DLL files on Windows 32-bit and Windows 64-bit platforms. The runtime DLL files on a 64-bit platform have the same base name as the 32-bit version with the addition of the "64" suffix. For example, `db2app.dll` on a Windows 32-bit operating system is equivalent to `db2app64.dll` on a Windows 64-bit operating system. For more information, see the *Call Level Interface Guide and Reference Volume 1*.

32-bit and 64-bit versions on Windows x64

DB2 does not support multiple DB2 32-bit and 64-bit versions installed on Windows, because the DB2 32 and 64-bit registries are stored in different locations. If you install the DB2 64-bit version, the 32-bit version is removed from the system.

LDAP and CLI configuration

With DB2 Version 8, if an application needs different LDAP settings, it must authenticate with a different LDAP user. Otherwise, the CLI configuration affects all DB2 copies that the LDAP user might potentially use.

Performance counters

Performance counters can be registered for only one DB2 copy at a time and they can monitor only the instances in the DB2 copy in which they are registered. When you switch the default DB2 copy, the DB2 Selection Wizard de-registers and

reregisters the performance counters so that they are active for the default DB2 copy.

Windows Management Instrumentation (WMI)

Only one version of the WMI provider can be registered at any given time.

Client Connectivity

You can use only one DB2 copy in the same process. For more information, see “IBM data server client connectivity using multiple copies” on page 135.

Applications that dynamically link DB2 DLL files

Applications that link to DB2 DLL files directly or that use LoadLibrary instead of LoadLibraryEx with the `LOAD_WITH_ALTERED_SEARCH_PATH` flag must ensure that the initial dependent library is loaded properly. You can use your own coding technique to check that the library loads, or you can call the `db2envvar.bat` file to set up the environment before running the application, or you can call the `db2SelectDB2Copy` API, which can be statically linked into the application.

Visual Studio plug-ins

If the default DB2 copy is a copy of Version 9.7 or earlier release, there can be only one version of the plug-ins registered on the same computer at the same time. The version of the plug-ins that is active is the version that is installed with the default DB2 copy.

Licensing

Licenses must be registered for each DB2 copy. They are not system-wide. Copy dependent licensing provides the ability for both restricted versions of DB2 products and full versions of DB2 products on the same host.

Windows services

DB2 services on Windows platforms use the `<servicename_installationname>` naming convention. For example, `DB2NETSECSERVER_MYCOPY1`. The display name also contains the Copy Name appended to it in brackets, for example, DB2 Security Server (*MYCOPY1*). Instances also include the *DB2-<DB2 Copy Name>-<Instance Name>-<Node Number>* in the display name, which is shown in the services control panel applet. The actual service name remains as is.

API to select the DB2 copy to use

You can use the `db2SelectDB2Copy` API to select the DB2 copy that you want your application to use. This API does not require any DLL files. It is statically linked into your application. You can delay the loading of DB2 libraries and call this API first before calling any other DB2 APIs.

Note: The `db2SelectDB2Copy` API cannot be called more than once for any given process; that is, you cannot switch a process from one DB2 copy to another.

The `db2SelectDB2Copy` API sets the environment required by your application to use the DB2 copy name or the location specified. If your environment is already set up for the copy of DB2 that you want to use, then you do not need to call this API.

If, however, you need to use a different DB2 copy, you must call this API before loading any DB2 DLL files within your process. This call can be made only once per process.

Database Partitioning with multiple physical nodes

Each physical partition must use the same DB2 copy name on all computers.

Using MSCS with Multiple DB2 Resources

Each DB2 resource must be configured to run in a separate resource monitor.

Multiple DB2 copies on the same computer (Linux and UNIX)

DB2 products can be installed to an installation path of your choice. DB2 products can also be installed multiple times on a single machine. Each DB2 Copy can be at the same or different code levels.

Root installations of DB2 products can be installed to an installation path of your choice. Unless specified otherwise, the default installation paths are:

- for AIX, HP-UX or Solaris operating systems /opt/IBM/db2/V10.1
- for Linux operating systems /opt/ibm/db2/V10.1

If you are using the DB2 Setup wizard and choose to install a new DB2 copy when the default installation path is already in use, the default path becomes:

- for AIX, HP-UX or Solaris /opt/IBM/db2/V10.1_##
- for Linux /opt/ibm/db2/V10.1_##

where ## is a zero-prefixed sequential number in the range 01 to 99. The first occurrence of modifying the default path is _01. For example, a DB2 product is installed using the DB2 Setup wizard without specifying an installation path, and, the default installation path /opt/ibm/db2/V10.1 is already in use, the DB2 product will be installed in the new default path /opt/ibm/db2/V10.1_01.

If you install a DB2 copy using an installation method other than the DB2 Setup wizard, you must specify the installation path. Automatic sequential incrementation of the default path only occurs when the DB2 Setup wizard is used to install a DB2 copy. Regardless of the installation method, another full product cannot be installed in the path of another DB2 copy (such as /opt/ibm/db2/V10.1). A DB2 copy installed to a subdirectory of an existing DB2 copy is not supported because all copies are rendered unusable.

With the ability to install multiple copies of DB2 products on your system and the flexibility to install DB2 products and features in the path of your choice, the **db2is** command helps you keep track of what is installed and where it is installed. Run the **db2is** command to locate DB2 products installed on your system.

The following restrictions apply when installing multiple DB2 copies on one system:

- Non-root installations do not support multiple DB2 copies
- There can be only one DB2 Administration Server (DAS) on the system.
- Instance names must be unique across DB2 copies.

- A DB2 copy is not aware of any instances created in another DB2 copy. However, the **db2iupdt** command can switch ownership of an instance from one DB2 copy to another.
- Creating links with the **db2ln** command for one DB2 copy renders the other copies nonfunctional. Links must not be created when multiple DB2 copies are intended to coexist.
- Linux 32-bit images cannot be installed on Linux x64 operating system platforms.
- The DB2 installation tools do not support multiple simultaneous invocations on the same machine by the same user.

IBM Secure Shell Server For Windows (for Optim Database Administrator usage)

The information in this document will help you get started with the IBM Secure Shell Server for Windows.

A new service called “IBM Secure Shell Server For Windows” will be installed together with the DB2 Server products. This service is a limited functionality port from the OpenSSH 5.4p1 package. This service provides a server service for the Optim™ Database Administrator, as well as other approved IBM products. This service is not for general usage outside this scope.

IBM Secure Shell Server For Windows uses the same configuration method as OpenSSH which is based on the `sshd_config` file. This configuration file can be found under `%ALLUSERSPROFILE%\Application Data\IBM\IBMSSH\etc` directory, where `%ALLUSERSPROFILE%` is the Windows provided environment variable which can be displayed using “set ALLUSERSPROFILE”. Not all server functionalities are provided, refer to the `sshd_config` file for the supported feature. Only the supported feature can be configured through this file.

Logs of this service can be viewed using the Windows Event Viewer, under the Application category. Search for Event Source - “IBM Secure Shell Server For Windows”.

To start or stop the IBM Secure Shell Server For Windows, you can use the Windows Services Microsoft Management Console. Alternatively, you can use the command interpreter to issue “net start IBM Secure Shell Server For Windows” or “net stop IBM Secure Server For Windows”.

Mounting DB2 product media

The procedure to mount your DB2 database product CD or DVD is different for AIX, HP-UX, Linux, and Solaris operating systems.

Mounting CDs or DVDs (AIX)

To mount your DB2 database product CD or DVD on AIX operating systems, use the System Management Interface Tool (SMIT).

Before you begin

Depending on your system configuration, you might need to log on with root user authority to mount discs.

Procedure

To mount the CD or DVD on AIX using SMIT, perform the following steps:

1. Insert the disc in the drive.
2. Create a disc mount point by entering the **mkdir -p /disc** command, where *disc* represents the CD or DVD mount point directory.
3. Allocate a disc file system using SMIT by entering the **smit storage** command.
4. After SMIT starts, select **File Systems > Add / Change / Show / Delete File Systems > CDROM File Systems > Add CDROM File System**.
5. In the Add a File System window:
 - a. Enter a device name for your CD or DVD file system in the **DEVICE Name** field. Device names for CD or DVD file systems must be unique. If there is a duplicate device name, you may need to delete a previously-defined CD or DVD file system or use another name for your directory. In this example, */dev/cd0* is the device name.
 - b. Enter the disc mount point directory in the **MOUNT POINT** window. In this example, the mount point directory is */disc*.
 - c. In the **Mount AUTOMATICALLY at system restart** field, select yes to enable automatic mounting of the file system.
 - d. Click **OK** to close the window, then click **Cancel** three times to exit SMIT.
6. Mount the CD or DVD file system by entering the **smit mountfs** command.
7. In the Mount a File System window:
 - a. Enter the device name for this CD or DVD file system in the **FILE SYSTEM name** field. In this example, the device name is */dev/cd0*.
 - b. Enter the disc mount point in the **Directory over which to mount** field. In this example, the mount point is */disc*.
 - c. Enter *cdarfs* in the **Type of Filesystem** field. To view the other kinds of file systems you can mount, click **List**.
 - d. In the **Mount as READ-ONLY system** field, select yes.
 - e. Accept the remaining default values and click **OK** to close the window.

Results

Your CD or DVD file system is now mounted. To view the contents of the CD or DVD, place the disk in the drive and enter the **cd /disc** command where *disc* is the disc mount point directory.

Mounting CDs or DVDs (HP-UX)

To mount your DB2 database product CD or DVD on HP-UX operating systems, issue the **mount** command.

Before you begin

Depending on your system configuration, you might need root user authority to mount discs.

Procedure

To mount your DB2 database product CD or DVD on HP-UX:

1. Insert the CD or DVD in the drive.

2. If necessary, define a new directory as the mount point for the CD or DVD drive. Define `/cdrom` as the mount point using the **`mkdir /cdrom`** command.
3. If necessary, identify the drive device file using the **`ioscan -fnC disk`** command. This command lists all recognized CD or DVD drives and their associated device files. The file name will be something similar to `/dev/dsk/c1t2d0`.
4. Mount the CD or DVD drive to the mount-point directory:

```
mount -F cdfs -o rr /dev/dsk/c1t2d0 /cdrom
```
5. Obtain a file listing to verify the mount using the **`ls /cdrom`** command.
6. Log out.

Results

Your CD or DVD file system is now mounted. View the contents of the CD or DVD by placing it in the drive and enter the **`cd /cdrom`** command where `cdrom` is the mount point directory.

Mounting the CD or DVD (Linux)

To mount a CD-ROM on Linux operating systems, issue the **`mount`** command.

Before you begin

Depending on your system configuration, you might need root user authority to mount discs.

Procedure

To mount the CD or DVD on Linux operating systems:

1. Insert the CD or DVD in the drive and enter the following command:

```
mount -t iso9660 -o ro /dev/cdrom /cdrom
```


where `/cdrom` represents the mount point of the CD or DVD.
2. Log out.

Results

Your CD or DVD file system is now mounted. View the contents of the CD or DVD by placing the disc in the drive and enter the **`cd /cdrom`** command where `cdrom` is the mount point directory.

Mounting CDs or DVDs (Solaris Operating Environment)

If the CD-ROM is not automatically mounted when you insert it into the drive on Solaris Operating System, issue the **`mount`** command.

Before you begin

If you are mounting the CD or DVD drive from a remote system using NFS, the CD or DVD file system on the remote computer must be exported with root access. Depending on your local system configuration, you might also need root access on the local computer.

Procedure

To mount the CD or DVD on Solaris:

1. Insert the CD or DVD into the drive.
2. If the Volume Manager (**vol**) is running on your system, the disc is automatically mounted as `/cdrom/cd_label` if the CD or DVD has a label or `/cdrom/unnamed_cdrom` if it is unlabeled.

If the Volume Manager is not running on your system, complete the following steps to mount the CD or DVD:

- a. Determine the name of the device by entering the following command:

```
ls -al /dev/sr* |awk '{print "/" $11}'
```

This command returns the name of the CD or DVD device. In this example, the command returns the string `/dev/dsk/c0t6d0s2`.

- b. Enter the following commands to mount the CD or DVD:

```
mkdir -p /cdrom/unnamed_cdrom  
mount -F hsfs -o ro /dev/dsk/c0t6d0s2 /cdrom/unnamed_cdrom
```

where `/dev/dsk/c0t6d0s2` represents the name of the device that was returned in the preceding step and `/cdrom/unnamed_cdrom` represents the CD or DVD mount directory.

3. Log out.

Results

Your CD or DVD file system is now mounted. View the contents of the CD or DVD by placing the disk in the drive and enter the **cd /cdrom** command where `cdrom` is the mount point directory.

Chapter 7. Installing DB2 database servers using the DB2 Setup wizard (Windows)

This task describes how to start the DB2 Setup wizard on Windows. Use the DB2 Setup wizard to define your installation and install your DB2 database product on your system.

Before you begin

Before you start the DB2 Setup wizard:

- If you are planning on setting up a partitioned database environment, refer to "Setting up a partitioned database environment".
- Ensure that your system meets installation, memory, and disk requirements.
- If you are planning to use LDAP to register the DB2 server in Windows operating systems Active Directory, extend the directory schema before you install, otherwise you must manually register the node and catalog the databases. For more information, see the "Extending the Active Directory Schema for LDAP directory services (Windows)" topic.
- You must have a local Administrator user account with the recommended user rights to perform the installation. In DB2 database servers where LocalSystem can be used as the DAS and DB2 instance user and you are not using the database partitioning feature, a non-administrator user with elevated privileges can perform the installation.

Note: If a non-Administrator user account is going to do the product installation, then the VS2010 runtime library must be installed before attempting to install a DB2 database product. The VS2010 runtime library is needed on the operating system before the DB2 database product can be installed. The VS2010 runtime library is available from the Microsoft runtime library download website. There are two choices: choose `vcredist_x86.exe` for 32-bit systems or `vcredist_x64.exe` for 64-bit systems.

- Although not mandatory, it is recommended that you close all programs so that the installation program can update any files on the computer without requiring a reboot.
- Installing DB2 products from a virtual drive or an unmapped network drive (such as `\\hostname\sharename` in Windows Explorer) is not supported. Before attempting to install DB2 products, you must map the network drive to a Windows drive letter (for example, Z:).

Restrictions

- You cannot have more than one instance of the DB2 Setup wizard running in any user account.
- The DB2 copy name and the instance name cannot start with a numeric value. The DB2 copy name is limited to 64 English characters consisting of the characters A-Z, a-z and 0-9.
- The DB2 copy name and the instance name must be unique among all DB2 copies.
- The use of XML features is restricted to a database that has only one database partition.

- No other DB2 database product can be installed in the same path if one of the following is already installed:
 - IBM Data Server Runtime Client
 - IBM Data Server Driver Package
 - *DB2 Information Center*
- The DB2 Setup wizard fields do not accept non-English characters.
- If you enable extended security on Windows Vista or Windows 2008, or higher, users must belong to the DB2ADMNS or DB2USERS group to run local DB2 commands and applications because of an extra security feature (User Access Control) that limits the privileges that local administrators have by default. If users do not belong to one of these groups, they will not have read access to local DB2 configuration or application data.

Procedure

To start the DB2 Setup wizard:

1. Log on to the system with the local Administrator account that you have defined for the DB2 installation.
2. If you have the DB2 database product DVD, insert it into the drive. If enabled, the autorun feature automatically starts the DB2 Setup Launchpad. If the autorun does not work, use Windows Explorer to browse the DB2 database product DVD and double-click the **setup** icon to start the DB2 Setup Launchpad.
3. If you downloaded the DB2 database product from Passport Advantage®, run the executable file to extract the DB2 database product installation files. Use Windows Explorer to browse the DB2 installation files and double-click the **setup** icon to start the DB2 Setup Launchpad.
4. From the DB2 Setup launchpad, you can view installation prerequisites and the release notes, or you can proceed directly to the installation. You might want to review the installation prerequisites and release notes for late-breaking information.
5. Click **Install a Product** and the Install a Product window displays the products available for installation.
 If you have no existing DB2 database products installed on your computer, launch the installation by clicking **Install New**. Proceed through the installation following the DB2 Setup wizard prompts.
 If you have at least one existing DB2 database product installed on your computer, you can:
 - Click **Install New** to create a new DB2 copy.
 - Click **Work with Existing** to update an existing DB2 copy, to add function to an existing DB2 copy, upgrade an existing DB2 Version 9.5 and Version 9.7 copy, or to install an add-on product.
6. The DB2 Setup wizard determines the system language, and launch the setup program for that language. Online help is available to guide you through the remaining steps. To invoke the online help, click **Help** or press **F1**. You can click **Cancel** at any time to end the installation.
7. Sample panels when using the DB2 setup wizard lead you to the installation process. See the related links.

Results

Your DB2 database product is installed, by default, in the *Program_Files\IBM\sqllib* directory, where *Program_Files* represents the location of the Program Files directory.

If you are installing on a system where this directory is already being used, the DB2 database product installation path has *_xx* added to it, where *xx* are digits, starting at 01 and increasing depending on how many DB2 copies you have installed.

You can also specify your own DB2 database product installation path.

What to do next

- Verify your installation.
- Perform the necessary post-installation tasks.

For information about errors encountered during installation, review the installation log file located in the *My Documents\DB2LOG* directory. The log file uses the following format: *DB2-ProductAbbrev-DateTime.log*, for example, *DB2-ESE-Tue Apr 04 17_04_45 2012.log*.

If this is a new DB2 product installation on Vista 64-bit, and you use a 32-bit OLE DB provider, you must manually register the *IBMDADB2* DLL. To register this DLL, run the following command:

```
c:\windows\SysWOW64\regsvr32 /s c:\Program_Files\IBM\SQLLIB\bin\ibmdadb2.dll
```

where *Program_Files* represents the location of the Program Files directory.

If you want your DB2 database product to have access to DB2 documentation either on your local computer or on another computer on your network, then you must install the *DB2 Information Center*. The *DB2 Information Center* contains documentation for the DB2 database system and DB2 related products. By default, DB2 information is accessed from the web if the *DB2 Information Center* is not locally installed.

IBM Data Studio can be installed by running the the DB2 Setup wizard

DB2 Express Edition and DB2 Workgroup Server Edition memory limits

If you are installing DB2 Express Edition, the maximum allowed memory for the instance is 4 GB.

If you are installing DB2 Workgroup Server Edition, the maximum allowed memory for the instance is 64 GB.

The amount of memory allocated to the instance is determined by the **INSTANCE_MEMORY** database manager configuration parameter.

Important notes when upgrading from Versions 9.5 or 9.7:

- The self tuning memory manager does not increase your overall instance memory limit beyond the license limits.

Chapter 8. Installing DB2 servers using the DB2 Setup wizard (Linux and UNIX)

This task describes how to start the DB2 Setup wizard on Linux and UNIX operating systems. The DB2 Setup wizard is used to define your installation preferences and to install your DB2 database product on your system.

Before you begin

Before you start the DB2 Setup wizard:

- If you are planning on setting up a partitioned database environment, refer to “Setting up a partitioned database environment” in *Installing DB2 Servers*
- Ensure that your system meets installation, memory, and disk requirements.
- Ensure you have a supported browser installed.
- You can install a DB2 database server using either root or non-root authority. For more information about non-root installation, see “Non-root installation overview (Linux and UNIX)” in *Installing DB2 Servers*.
- The DB2 database product image must be available. You can obtain a DB2 installation image either by purchasing a physical DB2 database product DVD, or by downloading an installation image from Passport Advantage.
- If you are installing a non-English version of a DB2 database product, you must have the appropriate National Language Packages.
- The DB2 Setup wizard is a graphical installer. You must have X windows software capable of rendering a graphical user interface for the DB2 Setup wizard to run on your machine. Ensure that the X windows server is running. Ensure that you have properly exported your display. For example, export `DISPLAY=9.26.163.144:0`.
- If you are using security software in your environment, you must manually create required DB2 users before you start the DB2 Setup wizard.

Restrictions

- You cannot have more than one instance of the DB2 Setup wizard running in any user account.
- The use of XML features is restricted to a database that is defined with the code set UTF-8 and has only one database partition.
- The DB2 Setup wizard fields do not accept non-English characters.
- For HP-UX 11i V2 on Itanium based HP Integrity Series Systems, users created with Setup Wizard for DB2 instance owner, fenced user, or DAS cannot be accessed with the password specified on DB2 Setup Wizard. After the setup wizard is finished, you need to reset the password of those users. This does not affect the instance or DAS creation with the setup wizard, therefore, you do not need to re-create the instance or DAS.

Procedure

To start the DB2 Setup wizard:

1. If you have a physical DB2 database product DVD, change to the directory where the DB2 database product DVD is mounted by entering the following command:

```
cd /dvdrom
```

where */dvdrom* represents the mount point of the DB2 database product DVD.

2. If you downloaded the DB2 database product image, you must extract and untar the product file.
 - a. Extract the product file:

```
gzip -d product.tar.gz
```

where *product* is the name of the product that you downloaded.

- b. Untar the product file:

On Linux operating systems

```
tar -xvf product.tar
```

On AIX, HP-UX, and Solaris operating systems

```
guntar -xvf product.tar
```

where *product* is the name of the product that you downloaded.

- c. Change directory:

```
cd ./product
```

where *product* is the name of the product that you downloaded.

Note: If you downloaded a National Language Package, untar it into the same directory. This will create the subdirectories (for example *./nlpack*) in the same directory, and allows the installer to automatically find the installation images without prompting.

3. Enter the **./db2setup** command from the directory where the database product image resides to start the DB2 Setup wizard.
4. The IBM DB2 Setup Launchpad opens. From this window, you can view installation prerequisites and the release notes, or you can proceed directly to the installation. You can also review the installation prerequisites and release notes for late-breaking information.
5. Click **Install a Product** and the **Install a Product** window will display the products available for installation.

Launch the installation by clicking **Install New**. Proceed through the installation following the DB2 Setup wizard's prompts.

6. Sample panels when using the DB2 setup wizard will lead you to the installation process. See the related links.

After you have initiated the installation, proceed through the DB2 Setup wizard installation panels and make your selections. Installation help is available to guide you through the remaining steps. To invoke the installation help, click **Help** or press F1. You can click **Cancel** at any time to end the installation.

Results

For non-root installations, DB2 database products are always installed in the *\$HOME/sql1ib* directory, where *\$HOME* represents the non-root user's home directory.

For root installations, DB2 database products are installed, by default, in one of the following directories:

AIX, HP-UX, and Solaris

```
/opt/IBM/db2/V10.1
```

Linux /opt/ibm/db2/V10.1

If you are installing on a system where this directory is already being used, the DB2 database product installation path will have `_xx` added to it, where `_xx` are digits, starting at 01 and increasing depending on how many DB2 copies you have installed.

You can also specify your own DB2 database product installation path.

DB2 installation paths have the following rules:

- Can include lowercase letters (a-z), uppercase letters (A-Z), and the underscore character (`_`)
- Cannot exceed 128 characters
- Cannot contain spaces
- Cannot contain non-English characters

The installation log files are:

- The DB2 setup log file. This file captures all DB2 installation information including errors.
 - For root installations, the DB2 setup log file name is `db2setup.log`.
 - For non-root installations, the DB2 setup log file name is `db2setup_username.log`, where *username* is the non-root user ID under which the installation was performed.
- The DB2 error log file. This file captures any error output that is returned by Java (for example, exceptions and trap information).
 - For root installations, the DB2 error log file name is `db2setup.err`.
 - For non-root installations, the DB2 error log file name is `db2setup_username.err`, where *username* is the non-root user ID under which the installation was performed.

By default, these log files are located in the `/tmp` directory. You can specify the location of the log files.

There is no longer a `db2setup.his` file. Instead, the DB2 installer saves a copy of the DB2 setup log file in the `DB2_DIR/install/logs/` directory, and renames it `db2install.history`. If the name already exists, then the DB2 installer renames it `db2install.history.xxxx`, where *xxxx* is 0000-9999, depending on the number of installations you have on that machine.

Each installation copy has a separate list of history files. If an installation copy is removed, the history files under this install path will be removed as well. This copying action is done near the end of the installation and if the program is stopped or aborted before completion, then the history file will not be created.

What to do next

- Verify your installation.
- Perform the necessary post-installation tasks.

IBM Data Studio can be installed by running the the DB2 Setup wizard

National Language Packs can also be installed by running the `./db2setup` command from the directory where the National Language Pack resides, after a DB2 database product has been installed.

On Linux x86, if you want your DB2 database product to have access to DB2 documentation either on your local computer or on another computer on your network, then you must install the *DB2 Information Center*. The *DB2 Information Center* contains documentation for the DB2 database system and DB2 related products.

DB2 Express Edition and DB2 Workgroup Server Edition memory limits

If you are installing DB2 Express Edition, the maximum allowed memory for the instance is 4 GB.

If you are installing DB2 Workgroup Server Edition, the maximum allowed memory for the instance is 64 GB.

The amount of memory allocated to the instance is determined by the **INSTANCE_MEMORY** database manager configuration parameter.

Important notes when upgrading from Versions 9.5 or 9.7:

- If the memory configuration for your Versions 9.5 or 9.7 DB2 database product exceeds the allowed limit, the DB2 database product might not start after upgrading to the current version.
- The self tuning memory manager will not increase your overall instance memory limit beyond the license limits.

Chapter 9. DB2 Setup wizard installation walkthrough

When installing DB2 Enterprise Server Edition using the DB2 Setup wizard, there are several panels. The following leads you through these panels.

Procedure

1. To start the IBM DB2 Setup Launchpad, run the **db2setup** command from the directory where the DB2 installation image resides. For example,

```
./db2setup -l /tmp/db2setup.log -t /tmp/db2setup.trc
```

The **-l** and **-t** recommended parameters enable DB2 logging and tracing. Where the IBM DB2 Setup Launchpad opens. From this window, you can view the DB2 Enterprise Server Edition release information, upgrade options and pre-installation planning, or you can proceed directly to the installation.

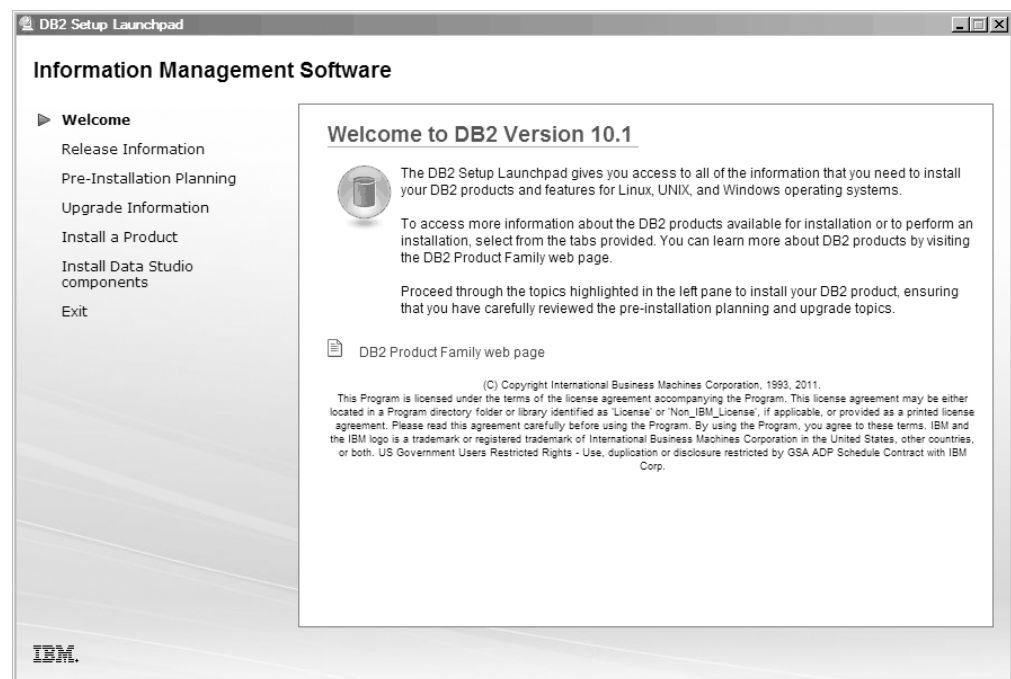


Figure 1. DB2 Setup Launchpad

2. To install a DB2 product, click **Install a Product**. The Install a Product panel displays. This window displays the products available for installation.



Figure 2. Install a Product window

Click **Install New**. The Welcome Panel displays.

3. The Welcome Panel. This panel displays the DB2 product you chose to install.




Figure 3. DB2 Setup wizard Welcome Panel

To end the installation, you can click you can click **Cancel** at any time.
Click **Next** to continue.

4. The Software License Agreement Panel. This panel displays the IBM and non-IBM licensing terms.




Figure 4. Software License Agreement Panel

After reviewing the software licensing agreement, select **Accept** to accept the terms.

Click **Next** to continue.

5. The Select the installation type Panel. On this panel, you specify the type of installation to perform. You can also view the features to be installed.

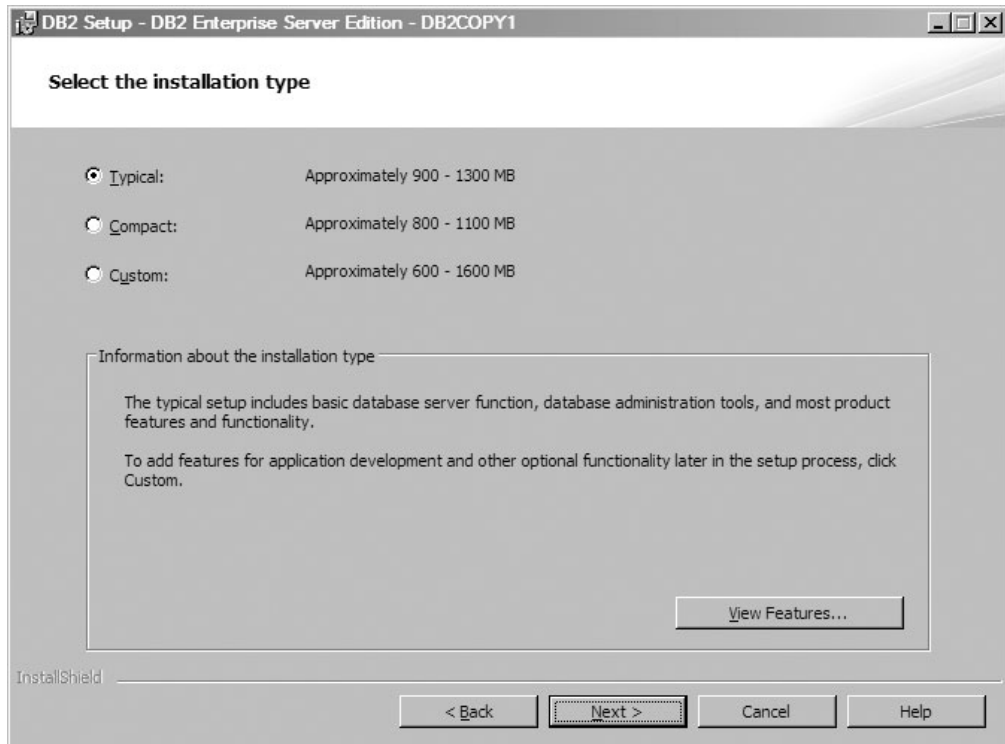


Figure 5. Select the installation type Panel

You can use the DB2 Setup wizard to specify :

- The typical setup includes basic database server function, database administration tools, and most product features and functionality. To add features for application development and other optional functionality later in the setup process, select **Custom**.
- The Compact setup includes basic DB2 features and functionality, and minimal configuration is performed.
- For Custom setup, if you are already familiar with DB2 features and settings, you can select features that you want installed and specify your configuration preferences. The default settings for a custom installation are the same as those for a typical installation, but you can customize optional features and configuration parameters.
- If you choose to setup using a default setting, you must specify the **Typical** field.

To view the features to be installed, click **View Features**. The View Features dialog box opens and displays the features to be installed.

Installation help can guide you through each of the installation panels. To invoke the installation help, click **Help** or press **F1**.

Click **Next** to continue.

6. The Select installation, response file creation, or both Panel. From this panel, you can choose to perform the installation now or, save the installation settings to a response file for installation at a later time or, perform both (perform the installation now and save the installation settings to a response file).




Figure 6. Select installation, response file creation, or both Panel

You can use the DB2 Setup wizard to:

- Install the product. When you are finished stepping through the DB2 Setup wizard, your DB2 database product is installed on your system. Response file is not saved.
- Create and save the installation settings in a response file, but not install the product. You can use this response file to install this product later. The DB2 database product is not installed. The DB2 Setup wizard only generates the response file based on your selections.
- Install the product and create and save a response file that you can use as a record of your choices.

If you choose to create a response file, you can specify the name and location of the response file in the **Response file name** field. This same response file can be updated to use on another set of hosts as well.

Click **Next** to continue.

7. The Select installation directory Panel. On this panel, you specify the directory where you want to install your DB2 database product.

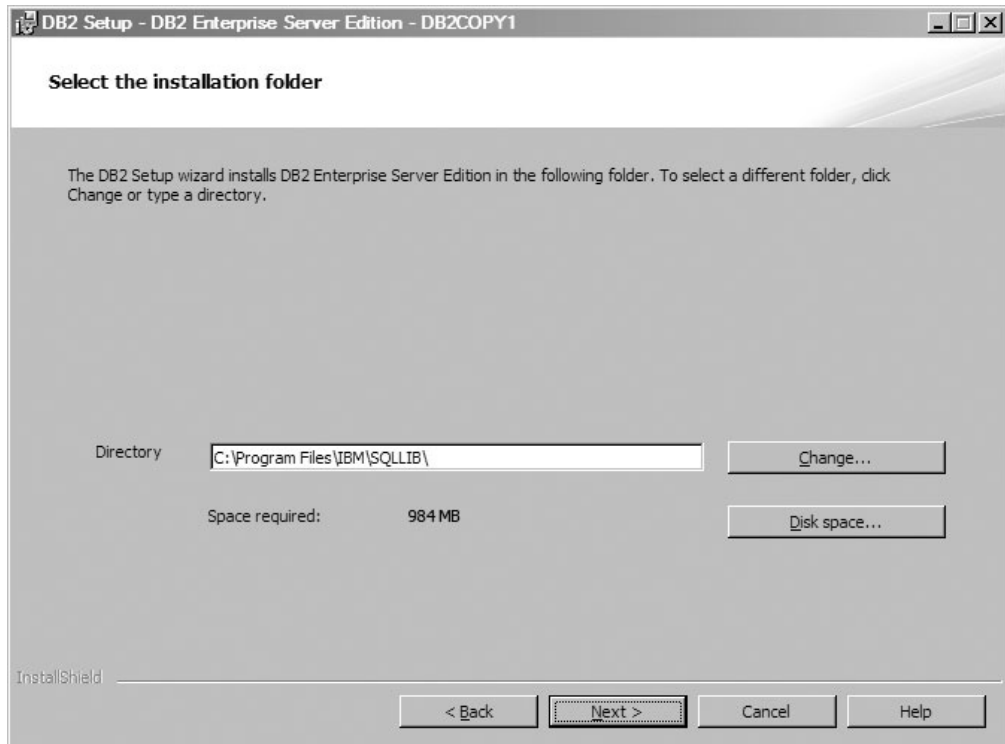


Figure 7. Select installation directory Panel

To view the disk space required for the installation of the selected features click **Disk space**. The Disk Space Requirement dialog box opens and displays the available disk space for the currently selected features.

Click **Next** to continue.

8. The Set user information for the default DB2 Administration Server (DAS). On this panel you to specify the domain, user name and password for the DB2 Administration Server.




Figure 8. Set user information for the default DB2 Administration Server

The DAS runs on your computer to provide support required by the DB2 tools. Specify the user information for the DAS.

Click **Next** to continue.

9. The Configure DB2 instances Panel. On this panel, you can view and change the DB2 instances.




Figure 9. Configure DB2 instances Panel

To view and change the DB2 instance configuration panel, click **Configure**. The TCP/IP port number is used by a DB2 instance to listen for incoming DB2 connections. The port number value must be in the range 1024 - 65535 and must be available on all hosts.

Click **Next** to continue.

10. The Set up notifications Panel. This panel you can manually configure your notifications to be sent to local or to remote contact list.




Figure 10. Set up notifications Panel

You can set up your DB2 server to automatically send notifications to alert administrators when a database needs attention. The contact information is stored in the administration contact list. You need an unauthenticated SMTP server to send these notifications. If you do not set up your DB2 server to send notification at this time, the health alerts are still recorded in the administration notification log.

You can set the administration contact option to be local (create a contact list on this computer) or remote (use an existing contact list on another DB2 server). You must specify the remote DB2 server.

Click **Next** to continue.

Note: The Set up notifications Panel is unavailable in the Typical Install GUI option for DB2 Express Edition and DB2 Express - C Edition. If you want this panel to be available during the installation process, you must choose the Custom Install GUI option.

11. The Specify a contact for health monitor notification Panel. On this panel, you can choose to create a new administration contact.




Figure 11. Specify a contact for health monitor notification Panel

By default, a health monitor runs on the DB2 instance you are setting up. You can either specify the administration contact for this instance or defer this task until after the installation is complete.

Click **Next** to continue.

Note: The Specify a contact for health monitor notification Panel is unavailable in the Typical Install GUI option for DB2 Express Edition and DB2 Express - C Edition. If you want this panel to be available during the installation process, you must choose the Custom Install GUI option.

12. The Enable operating system security for DB2 objects Panel. This panel enables operating system security.



Figure 12. Enable operating system security for DB2 objects Panel

To enable the operating system security, specify the administrator and user group information .

Click **Next** to continue.

Note: The Enable operating system security for DB2 objects Panel is unavailable in the Typical Install GUI option for DB2 Express Edition and DB2 Express - C Edition. If you want this panel to be available during the installation process, you must choose the Custom Install GUI option.

13. The Start copying files and create response file Panel. This panel outlines the information you need to create the response file. This is the final step before the installation wizard starts copying the program files.

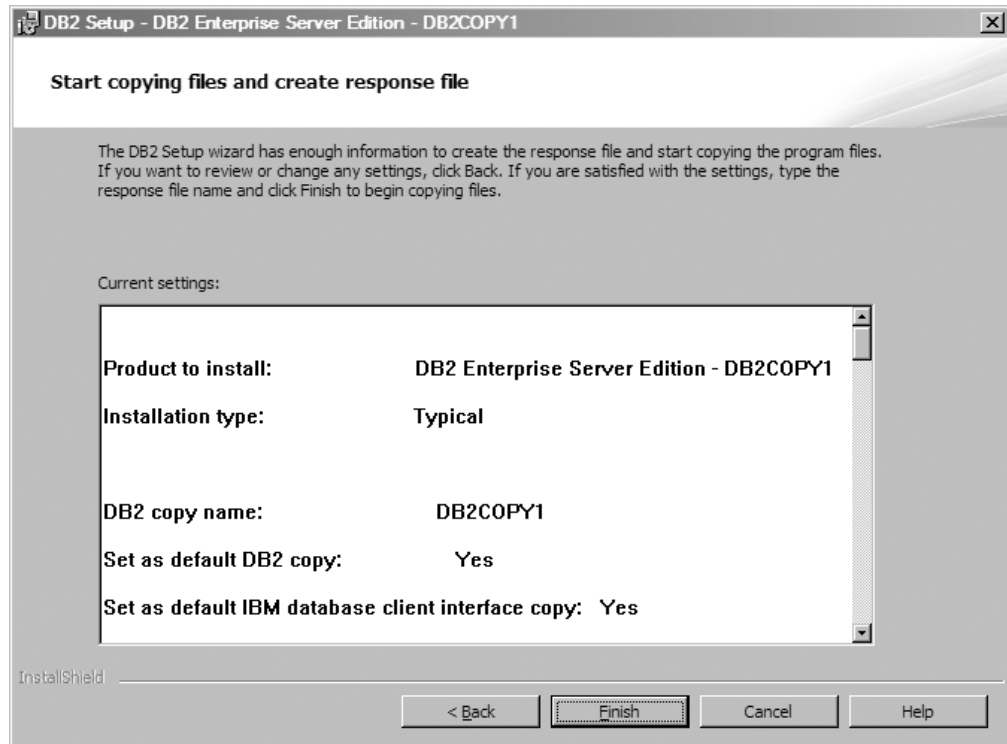


Figure 13. Start copying files and create response file Panel

This panel displays the choices necessary to complete the DB2 installation or to create a response file. Review all the features listed in this panel before clicking **Finish**. To change any of these settings, click **Back** to go back to the previous pages in the DB2 Setup wizard and change your choices. To launch the installation process across all specified hosts, click **Finish**.

14. The Setup is Complete Panel. This panel indicates the DB2 files are copied to your computer and the system configuration is complete.

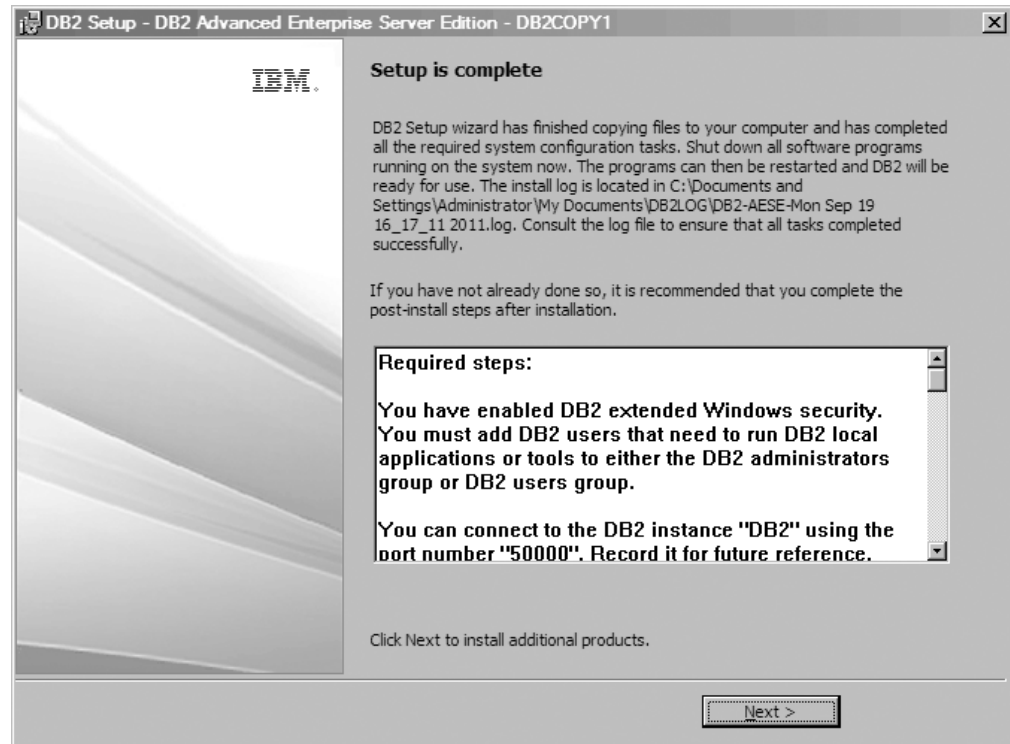


Figure 14. Setup is Complete Panel

The install log is located at C:\Documents and Settings\Administrator\My Documents. To use the DB2 database product, shut down all software programs running on the system and restarted the program.

15. The Install additional products Panel. This panel displays the additional products that can be installed after the DB2 database product installation.



Figure 15. Install additional products Panel

You can optionally install any of the following product by clicking the corresponding Install button. IBM Database Add-ins for Visual Studio provides application development support for .NET applications directly in Visual Studio 2008. To launch the installation process of IBM Database Add-ins for Visual Studio, click **Install**. After the installation is complete, you can review the post-installation steps, the log file for more details on what was set up, and the status of each remote host.

Click **Finish** to exit.

16. Optional. The Install Data Studio components Panel. Optionally, you can install the Data Studio components.




Figure 16. Install Data Studio components Panel

To install the Data Studio components, re-start the IBM DB2 Setup Launchpad by running the **db2setup** command. Click **Install Data Studio components**. The Install Data Studio components panel displays the products available for installation.

Chapter 10. Verifying the installation using the command line processor (CLP)

You can verify the installation by creating the SAMPLE database and running SQL commands to retrieve sample data.

Before you begin

- The SAMPLE database component, found in the features selection, must be installed on your system and is included in a typical installation.
- You require a user with SYSADM authority.

Procedure

To verify the installation:

1. Log on to the system as a user with SYSADM authority.
2. Start the database manager by entering the **db2start** command.
3. Enter the **db2samp1** command to create the SAMPLE database.

This command can take a few minutes to process. There is no completion message; when the command prompt returns, the process is complete.

The SAMPLE database is automatically cataloged with the database alias SAMPLE when it is created.

4. Connect to the SAMPLE database, retrieve a list of all the employees that work in department 20, and reset the database connection. Enter the following commands from the command line processor (CLP):

```
connect to sample
select * from staff where dept = 20
connect reset
```

The output should be similar to the following:

ID	NAME	DEPT	JOB	YEARS	SALARY	COMM
10	Sanders	20	Mgr	7	98357.50	-
20	Pernal	20	Sales	8	78171.25	612.45
80	James	20	Clerk	-	43504.60	128.20
190	Sneider	20	Clerk	8	34252.75	126.50

4 record(s) selected.

What to do next

After you have verified the installation, you can remove the SAMPLE database to free up disk space. Enter the **db2 drop database sample** command to drop the SAMPLE database.

Chapter 11. Taking the first steps after installation

After you install a DB2 database product, there are additional necessary and optional tasks available.

Deleting a Firefox browser profile

After you have run the DB2 Setup wizard or DB2 First Steps, you might want to delete the browser profile to prevent JavaScript from running automatically when visiting other Web pages with this profile.

About this task

If you use this profile to view other Web pages that have JavaScript, the JavaScript will run automatically without warnings.

If you use Firefox as your Web browser, a browser profile is created when you run DB2 First Steps. If you agree, a browser profile named DB2_FIRSTSTEPS is created.

Procedure

To delete a Firefox browser profile:

1. If any Firefox browsers are open, close them.
2. Open the Firefox Profile Manager.
 - On Windows
 - a. Open the Windows **Start** menu and select **Run...**
 - b. Type `firefox.exe -profilemanager` and click **OK**. The Firefox - Choose User Profile dialog opens.
 - On Linux and UNIX
 - a. Open a command prompt.
 - b. Switch to the Firefox directory.
 - c. Start Firefox with the **-profilemanager** switch. For example:
`./firefox -profilemanager`
3. Select the DB2_FIRSTSTEPS profile.
4. Click **Delete Profile...**
5. Click **Exit** to close the Firefox - Choose User Profile dialog.

What to do next

For further information about Firefox profiles, visit www.mozilla.org/support/firefox/profile

Applying DB2 licenses

To license your product, register the appropriate license key with your DB2 database product.

Before you begin

If you want the license key added automatically during product installation, copy the license key to the `/db2/license` directory of the installation image before launching the DB2 Setup wizard.

About this task

To use DB2 features or to be in compliance with your current product entitlement, you might be required to register additional license keys.

Example

For example, ensure you add the `db2ese_c.lic` or `db2ese_u.lic` license files to the `/db2/license` directory of the installation image before installing DB2 Enterprise Server Edition. If you do not add any license files to the `/db2/license` directory, you will see a “License not registered” message in the output from the **db2licm -l** command. You can apply the licenses manually after the installation by running the **db2licm -a** command.

If you register a valid license key, the **db2licm -l** command will successfully list all the products with available license information.

First Steps interface

Launches the First Steps interface, which contains links to the functions that you need to begin learning about and using the DB2 product.

From First Steps you can:

- Create the sample database or your own database.
- Check for product updates.

Product updates

Ensure that you are aware of the available product updates and enhancements for DB2 database products. With the update service you can view a list of available DB2 database product updates, and learn details about DB2 database product updates.

Running First Steps

You can manually start First Steps using the **db2fs** command:

```
$ db2fs
```

On Linux and UNIX operating systems

The **db2fs** command is located in the `INSTHOME/sqllib/bin` directory, where `INSTHOME` is the instance home directory. You need SYSADM authority to run this command.

Alternatively, on Linux operating systems you can start First Steps from **Main Menu** by selecting **IBM DB2 > First Steps**.

On Windows operating system

The **db2fs** command is located in the `DB2DIR\bin` directory, where `DB2DIR` is set to the location you specified during the DB2 installation.

Alternatively, you can start First Steps from the **Start** menu by selecting **Programs > IBM DB2 > [DB2 Copy Name] > Set Up Tools > First Steps** or from the DB2 database product folder on the desktop.

Post-installation tasks for DB2 database servers (Windows)

After installing DB2 database on Windows systems, you must complete a number of tasks.

Adding your user ID to the DB2ADMNS and DB2USERS user groups (Windows)

After successfully completing a DB2 installation, you now have to add users to the DB2ADMNS or the DB2USERS groups for users that need to run local DB2 applications and tools on the machine.

Before you begin

- You must have installed a DB2 database product.
- You must have selected the **Enable operating system security** check box on the Enable operating system security for DB2 object panel during the installation of your DB2 database product.

Procedure

To add users to the appropriate group:

1. Click **Start** and select **Run**.
2. Type **lusrmgr.msc** and click **OK**.
3. Select **Local Users and Groups**.
4. Select **Users**.
5. Select the user you want to add.
6. Click **Properties**.
7. Click the **Member Of** tab.
8. Click **Add**.
9. Select the appropriate group.
10. Click **OK**.

What to do next

If you did the install and chose not to enable the new security feature you can still do so post-install by running the **db2extsec.exe** command. Adding a user to a group takes effect the first time the user logs on after the user has been added. For example, if you add your user ID to the DB2ADMNS group, you need to log out and then log in again for this change to take effect.

Updating your 32-bit DB2 instances to 64-bit instances (Windows)

If you have more than one 32-bit DB2 copy on your system, you can convert them to 64-bit instances.

About this task

Restrictions

- DB2 32-bit and 64-bit instances cannot coexist on the same system.
- You cannot install a 64-bit DB2 copy on a system with multiple 32-bit DB2 copies.
- You cannot upgrade directly from a 32-bit pre-Version 10.1 DB2 copy to a 64-bit DB2 Version 10.1 copy. You can upgrade from a 32-bit pre-Version 10.1 DB2 copy to a 32-bit Version 10.1 DB2 copy then update from DB2 Version 10.1 32-bit to DB2 Version 10.1 64-bit. Refer to the "Upgrading DB2 32-bit servers to 64-bit systems" topic.

Procedure

1. Move all your 32-bit instances over to one DB2 copy using the **db2iupdt** command.
2. Uninstall the 32-bit DB2 copies, except for the one where you have moved all your 32-bit instances over to.
3. Install DB2 Version 10.1 64-bit which will upgrade your 32-bit DB2 Version 10.1 instances to 64-bit DB2 Version 10.1 instances.

Validating your DB2 copy

The **db2val** command ensures that your DB2 copy is functioning properly.

About this task

The **db2val** tool verifies the core function of a DB2 copy by validating installation files, instances, database creation, connections to that database, and the state of partitioned database environments. This validation can be helpful if you have manually deployed a DB2 copy on Linux and UNIX operating systems using tar.gz files. The **db2val** command can quickly ensure that all the configuration has been correctly done and ensure that the DB2 copy is what you expect it to be. You can specify instances and databases or you can run **db2val** against all of the instances. The **db2val** command can be found in the *DB2-install-path\bin* and *sqllib/bin* directories.

Example

For example, to validate all the instances for the DB2 copy, run the following command:

```
db2val -a
```

For complete **db2val** command details and further example, refer to the "db2val - DB2 copy validation tool command" topic.

Setting up notification and contact lists

For your DB2 database product to contact you or others regarding the status of your databases, set up notification and contact lists. If you did not do this during the DB2 database product installation, you can manually set up these lists.

Procedure

To set up notification and contact lists:

1. Log on as the instance owner or a user with SYSADM authority.

2. If the SMTP server was not specified during the installation process, you can configure it manually using the following command:

```
db2 update admin cfg using smtp_server host1
```

where host1 represents the TCP/IP host name for the SMTP server used for email notification.

3. If you want the contact information to be located on a remote DB2 administration server (DAS), you can specify the **contact_host** parameter using the following command:

```
db2 update admin cfg using contact_host host2
```

where host2 represents the TCP/IP host name where the DAS is running. If the **contact_host** parameter is not specified, the DAS assumes the contact information is local.

4. Turn on the scheduler using the following command:

```
db2 update admin cfg using sched_enable on
```

5. For these changes to take effect, restart the DAS use the following commands:

```
db2admin stop  
db2admin start
```

6. To make changes to the contact list, issue the **UPDATE HEALTH NOTIFICATION CONTACT LIST** command.

Changing the default DB2 and default IBM database client interface copy after installation (Windows)

The **db2swtch** command can be run from any DB2 copy or IBM data server driver copy.

Before you begin

Multiple DB2 copies (Version 9 or later) or multiple IBM database client interface copies (Version 9.5 or later) are installed on the same computer.

About this task

DB2COPY1

```
-ESE  
-WSE  
-...
```

DB2COPY1 is the default name of the DB2 copy that is the first installation of a DB2 database product on your machine. This same name is part of the instance directory where the database manager code and user data is stored.

If there are further DB2 copies installed on your machine, they receive default names DB2COPY2, DB2COPY3, and so on.

IBMDBCL1

DSDRIVER

IBMDBCL1 is the default name of the IBM Data Server Driver Package (DSDRIVER) that is the first installation of a driver on your machine.

If there are further DSDRIVER products installed on your machine, they receive default names: IBMDBCL2, IBMDBCL3, and so on.

Changing the default DB2 copy after installation (Windows)

To change the default DB2 copy using the Default DB2 and Database Client Interface Selection wizard:

1. Open the Default DB2 and Database Client Interface Selection wizard:
From a command window, run the **db2swtch** command without any additional parameters. Or, from the Start Menu, select **Programs > IBM DB2 > (DB2 copy name) > Set-up Tools > Default DB2 and Database Client Interface Selection wizard**. The Default DB2 and Database Client Interface Selection wizard opens.
2. On the Configuration page, select the copy type you want to work with (Default DB2 Copy).
3. On the default DB2 copy page, select the copy that you want to make the default and click **Next** to make it the default DB2 copy.
4. On the summary page, the wizard indicates the result of the operation.
5. Run the **dasupdt** command to move the DB2 Administration Server (DAS) to the new default copy.

This procedure switches the current default DB2 copy to the new default DB2 copy and makes the necessary changes to the registry. To access and use the new default DB2 copy, after you have moved the DAS to the new default DB2 copy, open a new command window. You can still access the original default DB2 copy by using the shortcuts in the Start menu for the original default DB2 copy.

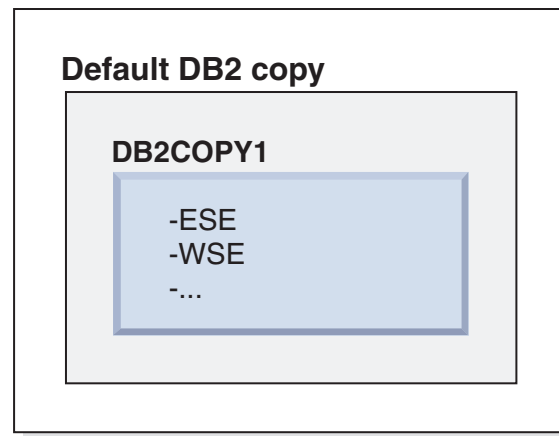
To change the default DB2 copy using the command line, you can use **db2swtch** command:

- The **db2swtch -d new_default_copy_name** command is used to change both the default DB2 copy and the default IBM database client interface copy to the new default copy name.
- The **db2swtch -db2 -d new_default_copy_name** command is used to change the default DB2 copy to the new default copy name.

This procedure removes registration of the current default DB2 copy and registers the specified DB2 copy as the default DB2 copy. It also makes the necessary changes to the registry, to the environment variables, to the ODBC and CLI drivers and the .NET data provider (when the change also makes the DB2 copy the default IBM database client interface copy), to the WMI registration, and to various other objects, and moves the DAS to the specified default DB2 copy. To access and use the new default DB2 copy, open a new command window.

Differences between Version 9.1 and Version 9.5 when changing the default DB2 copy In Version 9.1, after you install the first DB2 copy, it becomes the default DB2 copy.

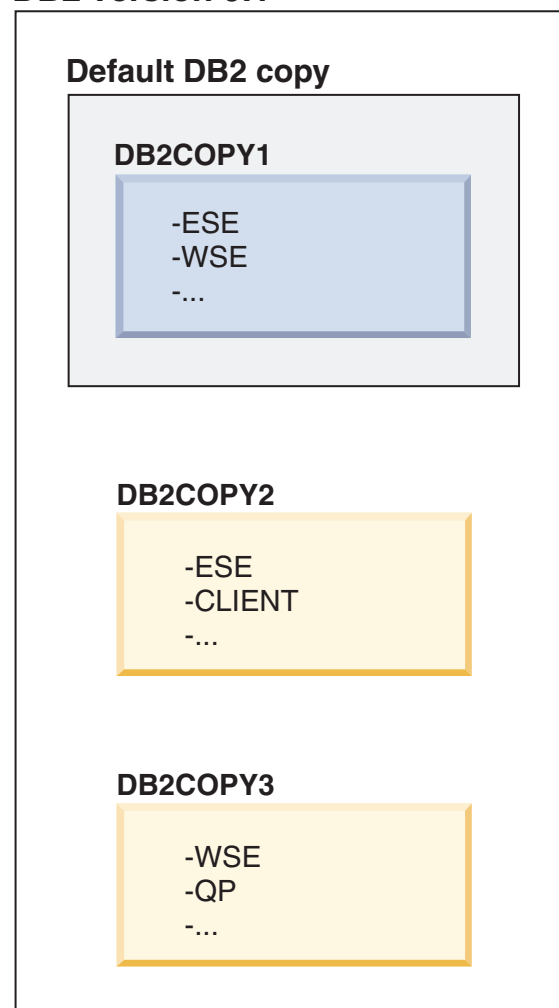
DB2 Version 9.1



Any database actions requiring the use of the DB2 server code access the code and the data that is from the default DB2 copy.

In Version 9.1, over time you can install other Version 9.1 or later DB2 copies.

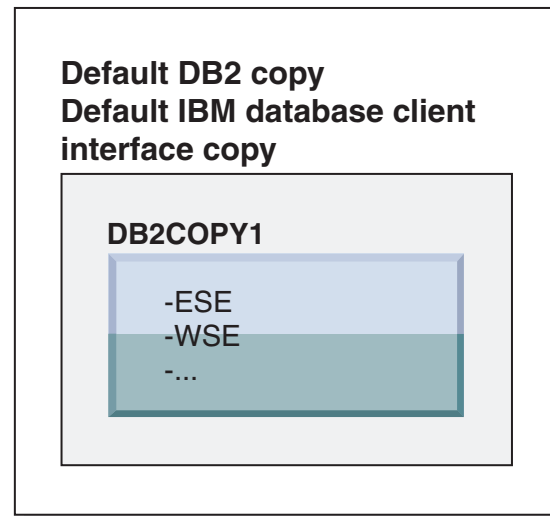
DB2 Version 9.1




However, only one DB2 copy is the default DB2 copy. If not explicitly referenced, requested database actions use the code and data that is from the default DB2 copy. If you want to use the code and data from the other DB2 copies (such as DB2COPY2 or DB2COPY3 in this example) and not from the default DB2 copy (DB2COPY1), then you must explicitly reference the code and data from the other DB2 copies.

In Version 9.5, after you install the first DB2 copy, it becomes the default DB2 copy and the default IBM database client interface copy.

DB2 Version 9.5




Legend




In Version 9.5, over time you can install other Version 9.5 or later DB2 copies.

DB2 Version 9.5


Default DB2 copy
Default IBM database client
interface copy




DB2COPY2



DB2COPY3



Legend



However, only one DB2 copy is the default DB2 copy or IBM Data Server driver copy. In this case, following the installation of two additional DB2 copies, both defaults remain associated with the original DB2 copy (DB2COPY1).

In Version 9.5, when you have installed more than one DB2 copy, you can choose to have one DB2 copy as the default DB2 copy and a different DB2 copy as the default IBM database client interface copy.

DB2 Version 9.5

Default DB2 copy

DB2COPY1

-ESE
-WSE
-...

Default IBM database client interface copy

DB2COPY2

-ESE
-CLIENT
-...

DB2COPY3

-WSE
-QP
-...

Legend



Default DB2 copy



Default IBM database
client interface copy

If not requiring the code associated with the default IBM database client interface copy (DB2COPY2), then applications or users requiring client interface code must explicitly reference the code that is from one of the other two DB2 copies (DB2COPY1 and DB2COPY3). If not explicitly referenced, the client interface code that is from the default IBM database client interface copy is used.

Changing the default IBM database client interface copy after installation (Windows)

To change the default IBM database client interface copy using the Default DB2 and Database Client Interface Selection wizard:

1. Open the Default DB2 and Database Client Interface Selection wizard:
From the Start Menu, select **Programs > IBM DB2 > (DB2 copy name) > Set-up Tools > Default DB2 and Database Client Interface Selection wizard**. Or go to the *installdir\bin* and run the **db2swtch** command.
The Default DB2 and Database Client Interface Selection wizard opens.
2. On the Configuration page, select the copy type you want to work with (Default IBM Database Client Interface Copy).
3. On the default IBM database client interface copy page, select the DB2 copy or the IBM data server driver copy that you want to make the default so that it is highlighted and click **Next** to make it the default IBM database client interface copy. (Recall that the code for the IBM database client interface copy was included with the DB2 copy.)
4. On the summary page, the wizard indicates the result of the operation.

This procedure switches the current default IBM database client interface copy to the new IBM data server driver copy making it the new default IBM database client interface copy. The procedure also makes the necessary changes to the registry. After the switching, the default ODBC and CLI drivers and the .NET data provider will point to the new copy. All the other ODBC drivers with copy name appended can also be accessed.

To change the default IBM database client interface copy using the command line, you can use **db2swtch** command:

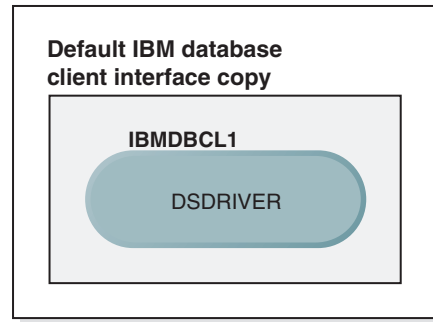
- The **db2swtch -d new_default_copy_name** command is used to change both the default DB2 copy and the default IBM database client interface copy to the new default copy name.
- The **db2swtch -client -d new_default_copy_name** command is used to change the default IBM database client interface copy to the new default copy name.

This procedure unregisters the current default IBM database client interface copy and registers the specified copy as the default IBM database client interface copy.


Changing the default IBM database client interface copy

In an environment with no other DB2 database products installed, you install the IBM Data Server Driver Package (DSDRIVER).

DB2 Version 9.5



Legend

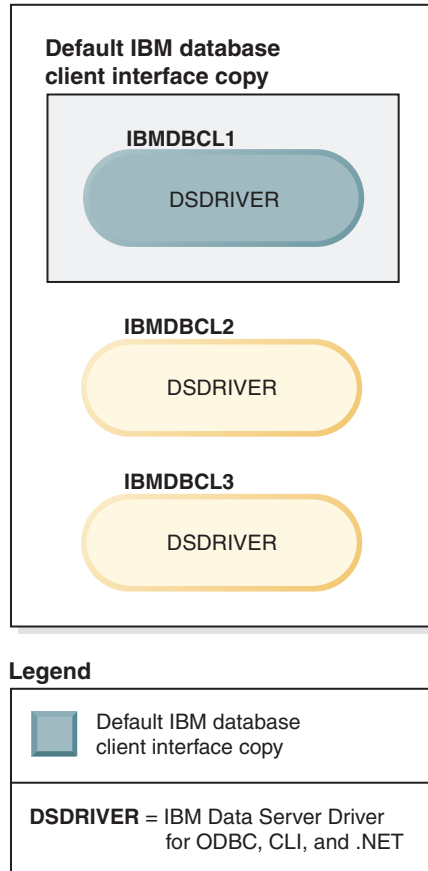
	Default IBM database client interface copy
DSDRIVER = IBM Data Server Driver for ODBC, CLI, and .NET	

By default, the name given to the DSDRIVER is IBMDBCL1.

With no other DB2 database server products (such as Enterprise Server Edition, or Workstation Server Edition) or other DSDRIVERS, this DSDRIVER is the default IBM database client interface copy. Any database actions requiring the use of application requester code access the code and data that is from the default IBM database client interface copy by default.

Over time, you can install DSDRIVERS in other IBM data server driver copies. However, only one IBM data server driver copy (or DB2 copy that is not shown in this example) is the default IBM database client interface copy.

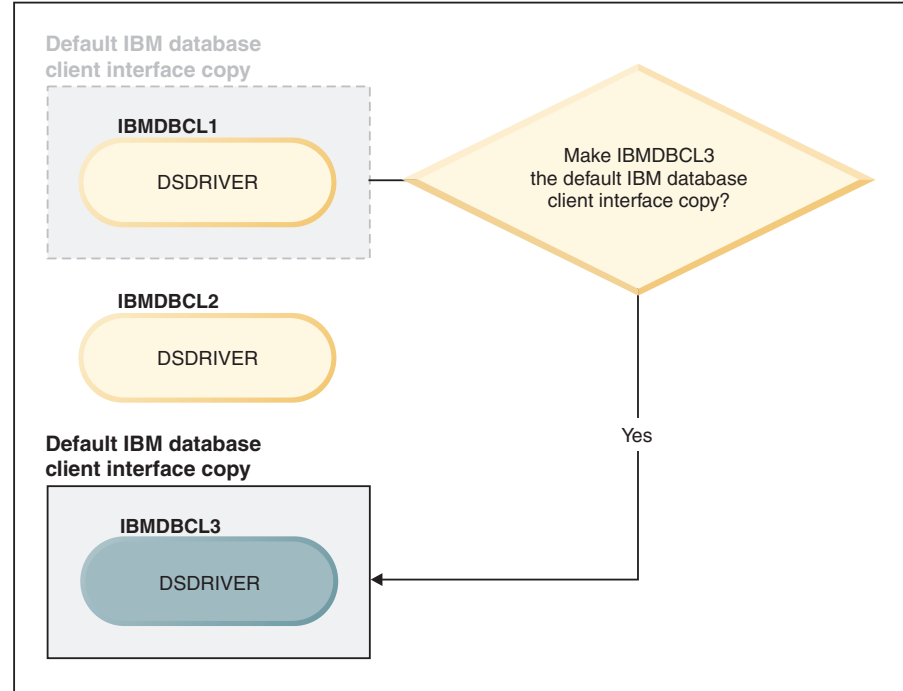
DB2 Version 9.5




If not explicitly referenced, application requests use the code and data that is from the default IBM database client interface copy. If you want to use the code from the other DSDRIVERS (such as IBMDBCL2 or IBMDBCL3 in the example) and not from the default IBM database client interface copy (IBMDBCL1), then you must explicitly reference the code and data from the other DSDRIVERS.

Over time, you can have several DSDRIVERS installed. Only one IBM database client interface copy is the default. At some point, you might decide to switch from one copy and make another copy the default IBM database client interface copy.

DB2 Version 9.5



Legend

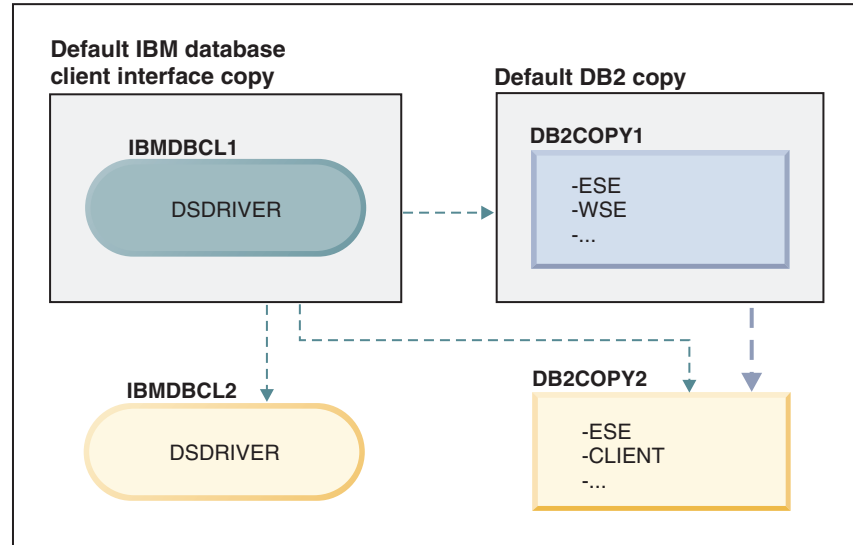
	Default IBM database client interface copy
DSDRIVER = IBM Data Server Driver for ODBC, CLI, and .NET	

Use the switch default DB2 copy and database client interface copy command (**db2swtch**) to choose and set the new default IBM database client interface copy. Using **db2swtch** with no arguments starts the Default DB2 and IBM Database Client Interface Selection wizard.

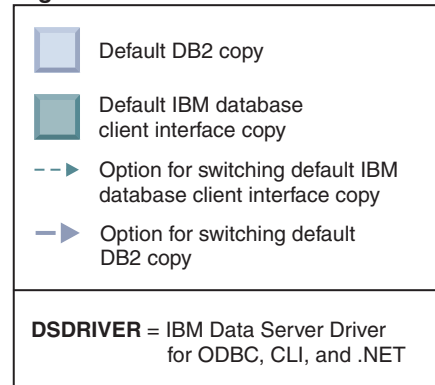
When both IBM database client interface copies and DB2 copies are on your machine

Over time you can have several DSDRIVERS and several DB2 copies installed on your machine. In the scenario presented here, IBMDBCL1 is the default IBM database client interface copy and DB2COPY1 is the default DB2 copy.

DB2 Version 9.5



Legend



If you want to change either or both defaults, use the **db2swtch** command with no arguments (Windows only) to start the Default DB2 and IBM Database Client Interface Selection wizard. The wizard shows all possible candidates when choosing a new default. You can also use the **db2swtch** command with **-client** or **-db2** option to do the switching.

When choosing a new default DB2 copy, in this scenario there is only one other choice: DBCOPY2.

When choosing a new default IBM database client interface copy, in this scenario there are three choices: IBMDBCL2, DB2COPY1, and DB2COPY2. (Recall that DB2 copies have the needed database client interface code to be declare the default IBM database client interface copy.)

IBM data server client connectivity using multiple copies

When using multiple DB2 copies of DB2 database products or multiple data server driver copies, various options are available for applications to access the DB2 databases. Existing applications will continue to function properly.

About this task

Restrictions

Only one copy can be used within the same process for each of the following modes of connecting to databases.

While this topic mainly deals with issues on Windows operating systems, on Linux and UNIX operating systems, you also cannot access instances associated with multiple DB2 copies within the same process.

Procedure

Review the considerations associated with each of the following methods to connect to DB2 databases:

- OLE DB

To use a DB2 copy other than the default, in the connection string, specify the IBM DADB driver name for this DB2 copy, which will be of the format: IBM DADB2.\$DB2_COPY_NAME. Some applications might not have the ability to change the connection strings without recompiling, therefore these applications will only work with the default DB2 copy. If an application uses the default program id, ibmdadb2, or the default clsid, it will always use the default DB2 copy.

Specifically, you will need to change the value of "provider=IBM DADB2" in the connection string. For example, if the DB2 copy that you want to use is called MY_COPY, you would specify "provider=IBM DADB2.MY_COPY" in the connection string. In case you must explicitly specify a GUID during installation, a response file keyword, **OLEDB_GUID**, is used to do this and allows you to enter your own GUID. Otherwise, the generated ID is used, as listed in the DB2 installation log.

Note: If you continue to use the IBM DADB2 provider name, then you will only be able to access data sources from the default DB2 copy.

- IBM Data Server Driver for ODBC and CLI

The IBM Data Server Driver for ODBC and CLI contains the copy name as part of the driver name. The default driver, IBM DB2 ODBC DRIVER, is set to the default IBM database client interface copy. The name of the driver for each installation is "IBM DB2 ODBC DRIVER - *Copy Name*".

Note:

- You can use only one copy in the same ODBC application at the same time
- Even when you set up a Data source with the default ODBC driver, it will be configured to access the DB2 copy that was the default at the time the Data source was cataloged.
- If you move or upgrade instances from one copy to another, you will re-configure the associated Data sources.

- IBM Data Server Provider for .NET

The IBM Data Server Provider for .NET is not accessed by the DB2 database product identified by the copy name. Instead, depending on the version of the provider that the application requires, it finds that version and uses it using the standard methods.

- JDBC or SQLJ

JDBC uses the current version of the driver in the classpath. By default, the classpath is configured to point to the default DB2 copy. Running **db2envvar.bat** from the DB2 copy you want to use will update your **PATH** and **CLASSPATH** settings for this copy.

- WMI

WMI does not support multiple DB2 copies. You can register only one copy of WMI at a time. To register WMI, follow this process:

- Unregister the WMI Schema extensions.
- Unregister the COM object.
- Register the new COM object.
- Use MOFCOMP to extend the WMI schema.

WMI is not registered during DB2 installation. You must complete the two registration steps. WMI is a selectable feature in DB2 database products. You must select this feature during a custom installation. It is not included in a typical installation.

- CLI applications

CLI applications that dynamically load the IBM data server client libraries should use the LoadLibraryEx API with the **LOAD_WITH_ALTERED_SEARCH_PATH** option, instead of the **LoadLibrary** option. If you do not use the LoadLibraryEx API with the **LOAD_WITH_ALTERED_SEARCH_PATH** option, specify db2app.dll in the **Path** by running **db2envvar.bat** from the bin directory of the DB2 copy that you want to use. For applications that link using db2api1.lib, to use a different DB2 copy, you can use the **/delayload** option in your **link** command to delay load db2app.dll and call the db2SelectDB2Copy API before any DB2 calls.

- DB2 System Tray

To reduce the number of system tray executables running on the system, by default any items in the system tray that are running in the previous default DB2 copy when the default copy is changed are disabled.

Working with existing DB2 copies

You can install and run multiple DB2 copies on the same computer. Each DB2 copy can either be at the same DB2 database product level or at a different DB2 database product level.

About this task

Restrictions

- On Linux and UNIX, non-root installations allow only one DB2 copy for each valid non-root user.
- On Windows, no other DB2 database product can be installed in the same path if either of the following products is already installed:
 - IBM Data Server Runtime Client
 - IBM Data Server Driver Package
 - *DB2 Information Center*

Procedure

1. When you start the DB2 Launchpad, it will display the available products that you can install.

You can install a new DB2 copy or work with an existing DB2 copy.

Note: Add-on products are always installed using the **Work with Existing** option.

2. Click **Work with Existing**.

The Work with Existing window opens and displays the existing DB2 copies on your system and the possible operations that can be performed. Possible operations include:

- Install
- Add new function
- Update (Windows only). Use this option to apply fix packs.
- Upgrade (Windows only). Use this option to upgrade your DB2 database product to a new release.
- Not available

Checking for DB2 updates

Check for product updates to ensure that you are aware of the available product updates and enhancements for DB2 products.

About this task

During DB2 product installation, the update service is enabled by default. The update service allows you to stay informed of product updates such as:

- Messages about releases and updates of the DB2 products.
- Availability of technical materials such as tutorials, Web cast, and white papers.
- IBM Marketing activities pertinent to your area of interest.

You can access product updates in any of the following ways:

- Using the command line
- Using First Steps
- On Linux operating systems, using the Main Menu.
- On Windows operating systems, using the shortcuts in the Start menu.

Restrictions

- This update service requires an Internet connection.
- On Windows operating systems, if the DB2 product was installed without elevated privileges, the update service is disabled.

Procedure

Access DB2 product updates and enhancements using one of the following ways:

- Using the command line, enter:
`db2updserv`
- Using First Steps:
 - You can start First Steps by typing the **db2fs** command.
 - On Windows operating systems, click **Start** and, select **Programs > IBM DB2 > [DB2 Copy Name] > Set-up Tools > First Steps**.

Click the **Start Check for Product Updates** button.

- On Windows operating systems, using the shortcuts in the Start menu, click **Start** and, select **Programs > IBM DB2 > [DB2 Copy Name] > Information > Check for DB2 Updates**.
- On Linux operating systems, click **Main Menu** and, select **IBM DB2 > Check for DB2 Updates**.

Results

With this update service you can view a list of available DB2 product updates, and learn details about DB2 product updates.

Installing the SSH Server component after installing the DB2 product (Windows)

The IBM Secure Shell (SSH) Server for Windows provides a secure platform to run commands from remote machines. If you chose not to install the IBM SSH Server for Windows component while installing a DB2 server, you can install this component manually. You require this component to use Optim tools.

Procedure

To install the IBM SSH Server for Windows:

1. From the Control Panel, select **Add or Remove Programs**.
2. Select the appropriate DB2 product, and click **Modify**. The DB2 Setup wizard opens.
3. On the Select the features to install panel, select the IBM Secure Shell Server for Windows component.
4. On the IBM Secure Shell Server panel, specify the SSH installation path and startup type.
5. Proceed through the DB2 Setup wizard panels, and click **Install**.

Post-installation tasks for DB2 database servers (Linux and UNIX)

After installing DB2 database on Linux or UNIX systems, you must complete a number of tasks.

Listing DB2 database products installed on your system (Linux and UNIX)

On supported Linux and UNIX operating systems, the **db21s** command lists the DB2 database products and features installed on your system, including the DB2 Version 10.1 HTML documentation.

Before you begin

At least one DB2 Version 9 (or later) database product must already be installed by a root user for a symbolic link to the **db21s** command to be available in the `/usr/local/bin` directory.

About this task

With the ability to install multiple copies of DB2 database products on your system and the flexibility to install DB2 database products and features in the path of your choice, you need a tool to help you keep track of what is installed and where it is installed. On supported Linux and UNIX operating systems, the **db21s** command lists the DB2 products and features installed on your system, including the DB2 HTML documentation.

The **db21s** command can be found both in the installation media and in a DB2 install copy on the system. The **db21s** command can be run from either location. The **db21s** command can be run from the installation media for all products except IBM Data Server Driver Package.

The **db21s** command can be used to list:

- Where DB2 database products are installed on your system and list the DB2 database product level

- All or specific DB2 database products and features in a particular installation path

Restrictions

The output that the **db2ls** command lists is different depending on the ID used:

- When the **db2ls** command is run with root authority, only root DB2 installations are queried.
- When the **db2ls** command is run with a non-root ID, root DB2 installations and the non-root installation owned by matching non-root ID are queried. DB2 installations owned by other non-root IDs are not queried.

The **db2ls** command is the only method to query a DB2 database product. You *cannot* query DB2 database products using Linux or UNIX operating system native utilities, such as **pkginfo**, **rpm**, **SMIT**, or **swlist**. Any existing scripts containing a native installation utility that you use to query and interface with DB2 installations must change.

You *cannot* use the **db2ls** command on Windows operating systems.

Procedure

- To list the path where DB2 database products are installed on your system and list the DB2 database product level, enter:

```
db2ls
```

The command lists the following information for each DB2 database product installed on your system:

- Installation path
- Level
- Fix pack
- Special Install Number. This column is used by IBM DB2 Support.
- Installation date. This column shows when the DB2 database product was last modified.
- Installer UID. This column shows the UID with which the DB2 database product was installed.
- To list information about DB2 database products or features in a particular installation path the **q** parameter must be specified:

```
db2ls -q -p -b baseInstallDirectory
```

where:

- **q** specifies that you are querying a product or feature. This parameter is mandatory.
- **p** specifies that the listing displays products rather than listing the features.
- **b** specifies the installation directory of the product or feature. This parameter is mandatory if you are not running the command from the installation directory.

Results

Depending on the parameters provided, the command lists the following information:

- Installation path. This is specified only once, not for each feature.
- The following information is displayed:

- Response file ID for the installed feature, or if the **p** option is specified, the response file ID for the installed product. For example, ENTERPRISE_SERVER_EDITION.
- Feature name, or if the **p** option is specified, product name.
- Product version, release, modification level, fix pack level (VRMF). For example, 10.1.0.0
- Fix pack, if applicable. For example, if Fix Pack 1 is installed, the value displayed is 1. This includes interim fix packs, such as Fix Pack 1a.
- If any of the product's VRMF information do not match, a warning message displays at the end of the output listing. The message suggests the fix pack to apply.

Validating your DB2 copy

The **db2val** command ensures that your DB2 copy is functioning properly.

About this task

The **db2val** tool verifies the core function of a DB2 copy by validating installation files, instances, database creation, connections to that database, and the state of partitioned database environments. This validation can be helpful if you have manually deployed a DB2 copy on Linux and UNIX operating systems using tar.gz files. The **db2val** command can quickly ensure that all the configuration has been correctly done and ensure that the DB2 copy is what you expect it to be. You can specify instances and databases or you can run **db2val** against all of the instances. The **db2val** command can be found in the *DB2-install-path\bin* and *sqllib/bin* directories.

Example

For example, to validate all the instances for the DB2 copy, run the following command:

```
db2val -a
```

For complete **db2val** command details and further example, refer to the “db2val - DB2 copy validation tool command” topic.

Main menu entries for DB2 tools (Linux)

After installation, you can add several DB2 tools to the Main menu.

On Linux operating systems, the following DB2 tools can be added to the Main menu:

- Check for DB2 Updates
- Command Line Processor (CLP)
- Command Line Processor Plus (CLPPlus)
- First Steps

These DB2 tools can be added to the main menu automatically or manually. Main menu entries are created automatically when any of the following DB2 commands are run:

- **db2icrt**
- **db2iupdt**
- **db2nrcfg**

- **db2nrupdt.**

The **db2icrt** and **db2iupdt** commands must be run by root. The **db2nrcfg** and **db2nrupdt** are for non-root install and are run by the instance owner.

To see the menu entries you might need to restart the desktop.

The main menu entries are automatically removed when either of the following commands are run:

- **db2_deinstall** (only removes the menu entries for the DB2 non-root instance when the non-root installation is removed)
- **db2idrop.**

In addition, you can manually create or remove main menu entries by running the following commands:

- **db2addicons** – to add menu entries
- **db2rmicons** – to remove menu entries.

Before running the **db2addicons** command, the DB2 instance environment must be set for the current user. The instance environment can be setup with:

Instance_HOME/sql1lib/db2profile (for Bourne shell and Korn shell users), or
Instance_HOME/sql1lib/db2chsrc (for C shell users), where *Instance_HOME* is the instance owner's home directory.

Part 3. Installing the DB2 pureScale Feature

If you are installing the IBM DB2 pureScale Feature then follow the “Installing the DB2 pureScale Feature” scenario.

If you have an existing DB2 database environment and want to upgrade to the DB2 pureScale Feature, then follow the converting instances scenario.

Before installing or upgrading to the DB2 pureScale Feature, it is critical that your preinstallation environment is properly set up and configured. After you have settled on an appropriate scenario that meets your needs, review and perform the preinstallation tasks to ensure that your system is properly set up.

If you want to remove the DB2 pureScale Feature, then follow the “Chapter 21, “Removing DB2 data server and the DB2 pureScale Feature,” on page 313” scenario.

Chapter 12. Network topology configuration considerations for DB2 pureScale environments

IBM DB2 pureScale Feature supports multiple communication adapter ports on the cluster caching facilities (CFs) and members.

Multiple communication adapter ports, enable DB2 support of network topologies with either one or two switches. A one-switch multiple communication adapter port configuration increases the throughput of requests. A two-switch configuration helps with increased throughput and high availability. DB2 pureScale environments do not require multiple communication adapter ports, however multiple communication adapter ports does help with resiliency and high availability.

Requirements for using multiple communication adapter ports

Multiple communication adapter ports are supported by DB2 pureScale Feature in the one-switch and two-switch configurations. The rules that apply to both configurations are as follows:

- Multiple communication adapter ports are supported on SLES, RHEL 6.1 and AIX on InfiniBand and 10GE networks.
- Members cannot be on the same host, or LPAR on AIX.
- The maximum number of communication adapter ports supported is four. The two validated and supported configurations for using multiple communication adapter ports are:
 - Four physical communication adapters, with one adapter port used by the CF or member on each adapter.
 - Two physical communication adapters, with two adapter ports on each adapter used by the CF or member.

Note: You can enhance high availability of adapter by using multiple physical communication adapters to connect to more than one switch. Using multiple communication adapter ports improves throughput.

- During installation and configuration, the cluster interconnect netnames you specify in the DB2 Setup wizard, or with the **db2icrt** and **db2iupdt** commands, are updated in the node configuration file, `db2nodes.cfg`. Host names that are not selected will not be listed in `db2nodes.cfg`.
- At least one switch is required in a DB2 pureScale environment.
- Two switches are required to support switch failover in a DB2 pureScale environment.
- IP subnets
 1. Each communication adapter port must be on a different subnetwork, also referred to as a subnet.
 2. If there are an equal number of communication adapter ports, each CF or member must be on the same set of subnets.
 3. If one CF server or member has fewer adapter ports than another, the one with more adapter ports must be on all the subnets as the CF or member with fewer adapter ports is on.
 4. If your members have only a single adapter, the communication adapter ports on all members must be on the same IP subnet. For simplicity, use the same the IP subnet of the first communication adapter port of the CF.

Members do not need to be on different IP subnets for availability reasons (adapter or switch failure) because the high speed communication between members and CFs through the switches uses different address resolution protocol than traditional interconnect (for example, Ethernet).

5. If you have multiple adapters on members and CFs, see Figure 2.
6. The netmask must be the same for all CFs and members.
7. Communication adapter ports that are used by applications other than DB2 applications must use a different subnet than any member or CF on the host.

One-switch configuration with multiple communication adapter ports

All CF and member hosts in a one-switch configuration are connected to the same switch by multiple communication adapter ports. The one-switch configuration is the simplest DB2 pureScale environment with multiple communication adapter ports to set up. The redundant communication adapter ports connecting each CF or member to the switch increases the bandwidth and the redundant adapter ports improve fault tolerance in the event one of the links fail. As long as a CF or member has one functional communication adapter port and a public Ethernet connection the CF or member remains operational. The following table is one-switch network topology example with multiple communication adapter ports to each CFs.

Table 18. Multiple communication adapter ports on CF hosts and members with 1 switch

Host	Cluster inter-connect adapter	Adapter port	Network interface name	Cluster interconnect netname	IP address	Subnetwork mask (Netmask)	Subnet
PrimaryCF	iba0	0	ib0	PrimaryCF-netname1	10.111.0.1	255.255.255.0	10.111.0.0
PrimaryCF	iba0	1	ib1	PrimaryCF-netname2	10.111.1.1	255.255.255.0	10.111.1.0
PrimaryCF	iba1	0	ib2	PrimaryCF-netname3	10.111.2.1	255.255.255.0	10.111.2.0
PrimaryCF	iba1	1	ib3	PrimaryCF-netname4	10.111.3.1	255.255.255.0	10.111.3.0
SecondaryCF	iba0	0	ib0	SecondaryCF-netname1	10.111.0.2	255.255.255.0	10.111.0.0
SecondaryCF	iba0	1	ib1	SecondaryCF-netname2	10.111.1.2	255.255.255.0	10.111.1.0
SecondaryCF	iba1	0	ib2	SecondaryCF-netname3	10.111.2.2	255.255.255.0	10.111.2.0
SecondaryCF	iba1	1	ib3	SecondaryCF-netname4	10.111.3.2	255.255.255.0	10.111.3.0
Member0	iba0	0	ib0	Member0-netname1	10.111.0.101	255.255.255.0	10.111.0.0
Member0	iba1	0	ib1	Member0-netname2	10.111.1.101	255.255.255.0	10.111.1.0
Member1	iba0	0	ib0	Member1-netname1	10.111.0.102	255.255.255.0	10.111.0.0
Member1	iba1	0	ib1	Member1-netname2	10.111.1.102	255.255.255.0	10.111.1.0
Member2	iba0	0	ib0	Member2-netname1	10.111.0.103	255.255.255.0	10.111.0.0
Member2	iba1	0	ib1	Member2-netname2	10.111.1.103	255.255.255.0	10.111.1.0
Member3	iba0	0	ib0	Member3-netname1	10.111.0.104	255.255.255.0	10.111.0.0
Member3	iba1	0	ib1	Member3-netname2	10.111.1.104	255.255.255.0	10.111.1.0

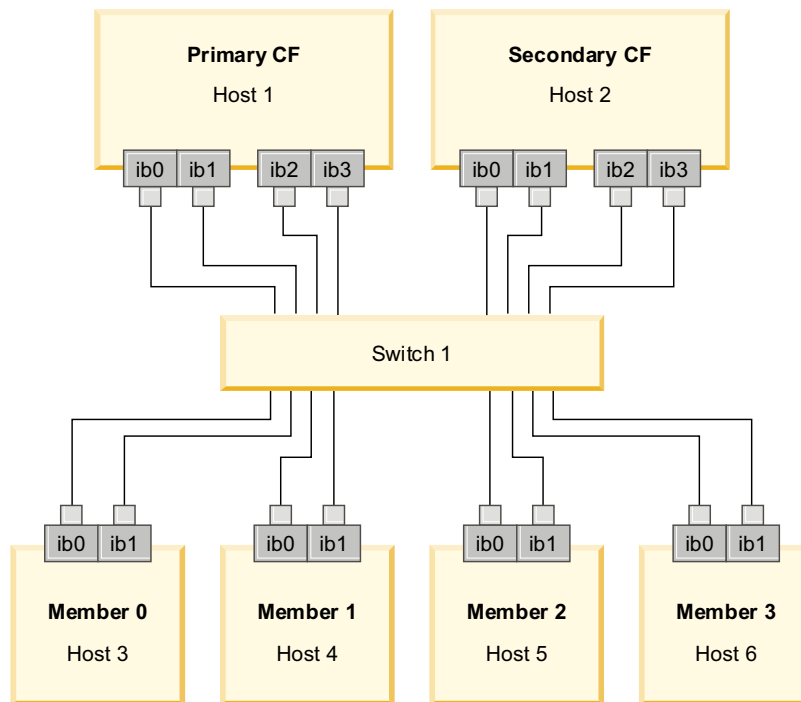


Figure 17. The 2 CFs and four members connect to one switch.

Note: Because IPoIB uses different address resolution protocol than other traditional interconnect (such as Ethernet), there is no requirement and benefit of putting members in different IP subnets in terms of availability of members in the case of adapter or switch failure.

Two-switch configuration with multiple communication adapter ports

The two-switch configuration improves network resiliency over the one-switch configuration. If a switch fails, the second switch reduces the risk of a critical failure. Additional setup requirements for a two-switch configuration are as follows:

- Half of the communication adapter ports must be connected to each switch.
- The switches must be connected to each other by two or more inter-switch links. Connect the two switches together by half the total number of cables that connect CFs and members, members to the switches to improve bandwidth and fault tolerance.
- Switch failover capability must be configured for the switch so that if one switch fails, the surviving switch and hosts connected to it are not impacted.
- Distribute the members evenly between the switches so that each switch is cabled to the same number of members.

Connect each CF or member adapter port with the same IP subnetwork to the same switch. If both ports of the same physical adapter are used, configure the adapter interface so that a port from the adapter is connected to each switch. By connecting to each switch by multiple adapters the risk posed by a switch or adapter failure is reduced:

- If an adapter of a CF or member fails, it can still communicate with each switch through the other surviving adapter, and a subsequent switch failure would not bring down the DB2 pureScale environment.

- If a switch fails, a subsequent adapter failure on a CF would still leave the primary and secondary CF intact.

The following table is two-switch network topology example with multiple communication adapter ports to each CF and member.

Table 19. Multiple communication adapter ports on CFs and members with two switches

Host	Cluster inter-connect adapter	Adapter port	Network interface name	Cluster interconnect netname	Connected to switch	IP address	Subnetwork mask (Netmask)	Subnet
PrimaryCF	iba0	0	ib0	PrimaryCF-netname1	1	10.222.0.1	255.255.255.0	10.222.0.0
PrimaryCF	iba0	1	ib1	PrimaryCF-netname2	2	10.222.1.1	255.255.255.0	10.222.1.0
PrimaryCF	iba1	0	ib2	PrimaryCF-netname3	1	10.222.2.1	255.255.255.0	10.222.2.0
PrimaryCF	iba1	1	ib3	PrimaryCF-netname4	2	10.222.3.1	255.255.255.0	10.222.3.0
SecondaryCF	iba0	0	ib0	SecondaryCF-netname1	1	10.222.0.2	255.255.255.0	10.222.0.0
SecondaryCF	iba0	1	ib1	SecondaryCF-netname2	2	10.222.1.2	255.255.255.0	10.222.1.0
SecondaryCF	iba1	0	ib2	SecondaryCF-netname3	1	10.222.2.2	255.255.255.0	10.222.2.0
SecondaryCF	iba1	1	ib3	SecondaryCF-netname4	2	10.222.3.2	255.255.255.0	10.222.3.0
Member0	iba0	0	ib0	Member0-netname1	1	10.222.0.101	255.255.255.0	10.222.0.0
Member0	iba0	1	ib1	Member0-netname2	2	10.222.1.101	255.255.255.0	10.222.1.0
Member1	iba0	0	ib0	Member1-netname1	1	10.222.0.102	255.255.255.0	10.222.0.0
Member1	iba0	1	ib1	Member1-netname2	2	10.222.1.102	255.255.255.0	10.222.1.0
Member2	iba0	0	ib0	Member2-netname1	1	10.222.0.103	255.255.255.0	10.222.0.0
Member2	iba0	1	ib1	Member2-netname2	2	10.222.1.103	255.255.255.0	10.222.1.0
Member3	iba0	0	ib0	Member3-netname1	1	10.222.0.104	255.255.255.0	10.222.0.0
Member3	iba0	1	ib1	Member3-netname2	2	10.222.1.104	255.255.255.0	10.222.1.0

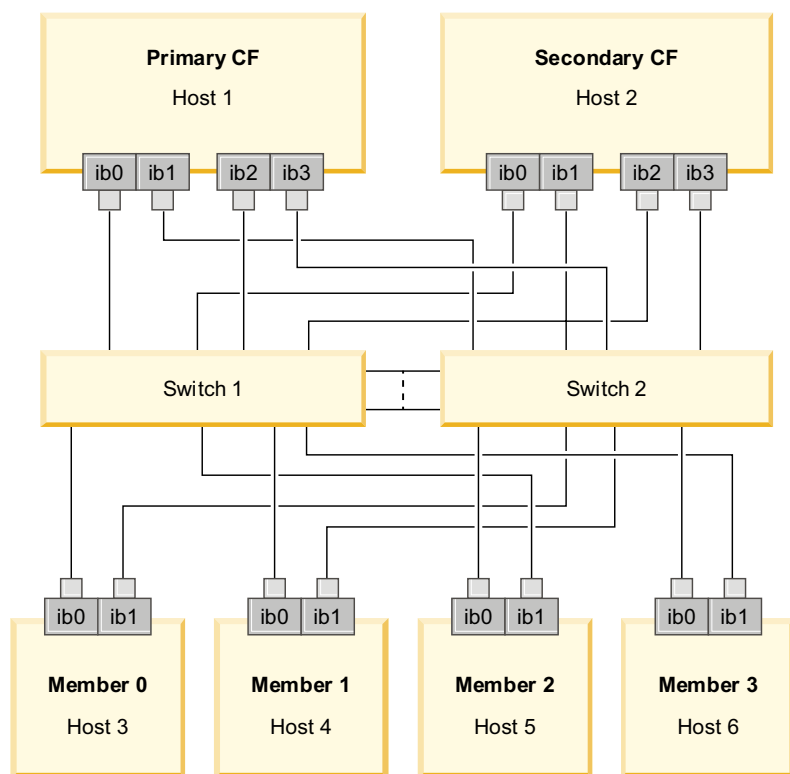


Figure 18. The 2 CFs and four members connect to two switches.

Configurations without multiple communication adapter ports

The following section is for illustration purposes. Configurations without multiple communication adapter ports do not offer redundancy on the switch.

In DB2 pureScale environments without multiple communication adapter ports, all member and CF communication adapter ports must be on the same subnet. As additional members are added, more CF resources are required to handle the members requests. If the number or amount of time members wait for CFs as members are added start to affect service level agreements of applications, consider adopting a multiple communication adapter ports topology.

The following table contains an example network topology of a DB2 pureScale environment with a single communication adapter port to each CF and member.

Table 20. Single communication adapter port with 1 switch

Host	Cluster interconnect adapter	Adap-ter port	Network interface name	Cluster interconnect netname	IP address	Subnetwork mask (Netmask)	Subnet
PrimaryCF	iba0	0	ib0	PrimaryCF-netname1	10.123.0.1	255.255.255.0	10.123.0.0
SecondaryCF	iba0	0	ib0	SecondaryCF-netname1	10.123.0.2	255.255.255.0	10.123.0.0
Member0	iba0	0	ib0	Member0-netname	10.123.0.101	255.255.255.0	10.123.0.0
Member1	iba0	0	ib0	Member1-netname	10.123.0.102	255.255.255.0	10.123.0.0
Member2	iba0	0	ib0	Member2-netname	10.123.0.103	255.255.255.0	10.123.0.0
Member3	iba0	0	ib0	Member3-netname	10.123.0.104	255.255.255.0	10.123.0.0

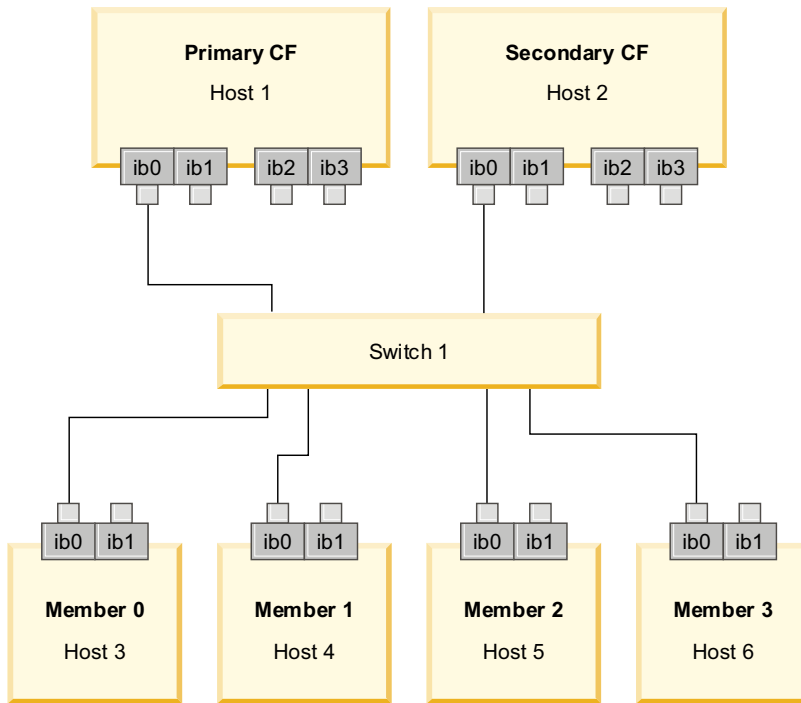


Figure 19. The 2 CFs and four members connect to one switch.

Chapter 13. Shared storage considerations for DB2 pureScale environments

Shared storage devices supported by IBM DB2 pureScale Feature are divided into three categories. These categories group storage devices and multipath I/O driver pairs based on two storage features.

The two features, fast I/O fencing support, and DB2 cluster services tiebreaker support, provide faster recovery times for failed hosts and better resiliency. In a production environment, these storage features improve availability and performance is also improved.

DB2 pureScale Feature supports all storage area network (SAN) and directly attached shared block storage, referenced as a logical unit number (LUN).

Fast I/O fencing

Before recovering a failed member in the DB2 pureScale instance, DB2 cluster services ensures that it is not possible for the failed member to modify shared data on disk. This precaution is known as I/O fencing. Some storage controllers and multipath I/O driver combinations support specialized SCSI commands, known as SCSI-3 Persistent Reserve (PR), with the write exclusive all registrants persistent reservation type. The SCSI-3 PR type write exclusive all registrants (type 7h) enables DB2 cluster services to perform fast I/O fencing, in as little as 1 - 2 seconds.

Recovery time is the elapsed time from a host failure until the release of locks for uncommitted transactions. For a typical OLTP workload with short transactions, fast I/O fencing results in recovery times of approximately 20 seconds. Fast I/O fencing is much faster than the alternative method, based on expiry of lock leases. When deciding which category to use, evaluate if fast I/O fencing support is required to meet your service level agreements (SLAs).

DB2 cluster services tiebreaker support

A subcluster of hosts requires operational quorum for IBM Reliable Scalable Cluster Technology (RSCT) to allow the subcluster to continue running the DB2 pureScale instance. To be granted operational quorum by RSCT, a subcluster must contain a majority of the hosts. If exactly half of the hosts in cluster remain in communication, the tiebreaker can be exclusively reserved by that subcluster to obtain operational quorum.

At the time of installation, the DB2 installer verifies whether a DB2 cluster services tiebreaker is specified and whether it can be used. Installation and instance creation is successful without a tiebreaker, however its absence results in a warning message that indicates a tiebreaker was not configured. Consider deploying an odd number of hosts to reduce the risk that failures leave a subcluster of half the original hosts.

On supported Linux systems, the device used as the cluster manager (RSCT) tiebreaker disk must have SCSI-3 PR WRITE EXCLUSIVE REGISTRANTS ONLY (reservation type code 0x50) enabled.

On supported AIX operating systems, this tiebreaker disk is not a requirement because the tiebreaker reservation is using SCSI-2 reserve/release mechanism.

Protection from the case where precisely half the hosts fail is of particular importance in a two machine configuration with an even number of hosts. In this configuration, where a single machine failure causes half the hosts to fail, a tiebreaker is needed to obtain operational quorum.

Category 1 storage device and multipath I/O driver combinations

Storage devices and multipath I/O driver combinations listed in this category can successfully support both the DB2 cluster services tiebreaker and fast I/O fencing. Category 1 devices have been validated with the DB2 pureScale Feature and result in the highest resiliency and fastest recovery times.

Table 21. Category 1 storage device and multipath I/O driver combinations

Storage Devices	Multipath I/O drivers required for AIX systems	Multipath I/O drivers required for Linux systems	Protocol
IBM System Storage® DS8000® series	SDDPCM	DM-MP	Fibre Channel
IBM System Storage DS5000 series	MPIO	DM-MP or RDAC	Fibre Channel
IBM System Storage DS4000® series	MPIO	DM-MP or RDAC	Fibre Channel
IBM System Storage DS3000 series	MPIO	DM-MP or RDAC	Fibre Channel
EMC VMAX/Symmetrix family ¹	MPIO driver provided by EMC (driver file EMC.Symmetrix.fcp.MPIO.rte)	DM-MP	Fibre Channel
NetApp FAS filers	MPIO driver provided by NetApp	DM-MP	iSCSI
Virtual I/O Server (VIOS) ²	MPIO or SDDPCM		Fibre Channel
Hitachi Virtual Storage Platform (VSP) ⁵	MPIO driver provided by IBM or HDLM		Fibre Channel

Note:

1. Each host port connected to the EMC VMAX/Symmetrix family storage system must have the SCSI3 (SC3) flag set on the front-end adapter port configuration. On the hypervolume, the logical unit number (LUN), used for the DB2 cluster services tiebreaker the SCSI3_persist_reserve option must be set.

On supported Linux systems, the device selected for cluster manager (RSCT) tiebreaker, must support reservation type of WRITE EXCLUSIVE REGISTRANTS ONLY type code 5h. With SCSI-3 PR, there is one persistent reservation holder, which is the node that acquires the quorum device. A DB2 pureScale cluster using EMC VMAX shared storage, must have the SCSI3_persist_reserve type option enabled for the disks mapped. It is not enabled by default.

2. A Virtual I/O Server must be used with N_Port ID Virtualization (NPIV) for a storage system and an AIX system multipath I/O driver combination. The VIOS storage is in the same category of storage that the underlying physical hardware storage system is in. For VIOS storage to be considered category 1 storage the physical storage hardware and driver combination must be in the category 1 table.

3. Starting from SDDPCM v2630 and higher, before installing the DB2 pureScale Feature, you must change the timeout_policy attribute to retry_path on the tiebreaker disk. To change the value of the timeout_policy attribute, enter the following command:

```
chdev -l hdiskX -a timeout_policy=retry_path
```

For details, see: http://www-01.ibm.com/support/docview.wss?rs=540&context=ST52G7&dc=D600&uid=ssg1S1004072&loc=en_US&cs=utf-8&lang=en

4. For more information about these drivers, see:

- For RDAC: <http://www.lsi.com/sep/Pages/rdac/index.aspx>
- For DM-MP on RHEL 5.5: http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/5/html/DM_Multipath/index.html
- For DM-MP on RHEL 6.1: http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/DM_Multipath/index.html
- For DM-MP on SUSE: http://www.novell.com/documentation/sles11/stor_admin/?page=/documentation/sles11/stor_admin/data/bookinfo.html

5. Minimum micro code level 70-04-31-00/00. Host Option Mode 72 must be on. HDLM for AIX Version 7.3.1 or later provided by Hitachi.

Category 2 storage device and multipath I/O driver combinations

Storage devices and multipath I/O driver combinations listed in this category are supported and have been validated with the DB2 pureScale Feature. The combinations in this category can successfully support the DB2 cluster services tiebreaker, but not fast I/O fencing.

Table 22. Category 2 storage device and multipath I/O driver combinations

Storage Devices	Multipath I/O drivers required for AIX systems	Multipath I/O drivers required for Linux systems	Protocol
IBM Storwize® V7000	SDDPCM or MPIO	DM-MP	Fibre Channel
IBM SAN Volume Controller	SDDPCM	DM-MP	Fibre Channel
EMC VMAX/Symmetrix family	EMC PowerPath		Fibre Channel
Hitachi Universal Storage Platform V (USP V)	MPIO		Fibre Channel

Category 3 storage device and multipath I/O driver combinations

All storage devices and multipath I/O driver combinations not identified as category 1 or category 2 fall into category 3. Category 3 combinations have not been validated with the DB2 pureScale Feature, though they are supported.

Recovery times for host failures in a DB2 pureScale environment are competitive with other vendors solutions, regardless of the category of storage device and multipath I/O driver.

Chapter 14. User-managed file system

While the DB2 managed file system is recommended, you might want to use your existing shared file system.

Currently, the only supported user-managed file system is an IBM General Parallel File System (GPFS™) file system. For the supported version and level of the existing GPFS, see the installation prerequisites. During the product installation, the DB2 installer will check if the user-managed GPFS cluster and file system meet this requirement. If not, instance creation will fail.

In addition to the correct version level, you must verify your existing GPFS file system configuration. To verify the configuration of the current file system cluster, run the following command:

```
db2cluster -cfs -verify -configuration
```

To verify the file system setup, run the following command:

```
db2cluster -cfs -verify -filesystem mygpfs1
```

where mygpfs1 is the name of the existing GPFS cluster file system.

A DB2 Managed File System and a user-managed file system cannot coexist. IBM DB2 pureScale Feature supports only one DB2 pureScale instance; either you use a DB2 Managed File System or a user-managed file system.

If using a user-managed file system, the default password-less remote root login is required.

When you create your first DB2 pureScale instance, you must provide the full path to an existing GPFS file system to indicate that you want to use a user-managed file system. The DB2 installer will verify that the file system is valid to use.

For response file installations:

- the `INSTANCE_SHARED_DEVICE_PATH` keyword represents a DB2 Managed File System
- the `INSTANCE_SHARED_DIR` keyword represents a user-managed file system

For a DB2 Setup wizard installation, the **DB2 cluster file system** panel will detect the cluster environment.

For the **db2icrt** or **db2iupdt** command:

- the `-instance_shared_dev` option represents a DB2 Managed File System
- the `-instance_shared_dir` option represents a user-managed file system

For example:

```
db2icrt -d
        -cf host1:host1-ib0
        -m host2:host2-ib0
        -instance_shared_dir mygpfs1
        -tbdev /dev/hdisk2
        -u db2sdfel
        db2insd1
```

For this example, make sure the following are ready before running the **db2icrt** command:

- The mygpfs1 file system is the mount point of the existing GPFS cluster file system
- The same GPFS cluster is setup on both host1 and host2, and, the mygpfs1 file system is mounted on both of them
- The existing instance owner user db2insd1 has write access to the mygpfs1 directory from both host1 and host2
- Since this example is using a user-managed file system, specify `-instance_shared_dir` rather than `-instance_shared_dev`.

When extending the DB2 pureScale instance to a remote host, extend the cluster to the remote host and set up the host to have write access as root to the file system first. The DB2 installer will perform the corresponding validation.

For example:

```
db2iupdt -d -add -m host3:host3-ib0 db2insd1
```

In this example, running this command extends the instance to a new member (host3). Before running the **db2iupdt -add** command:

- host3 has to be included in the same GPFS cluster (which already has host1 and host2)
- the mygpfs1 file system has to be mounted from host3
- The existing instance owner user db2insd1 has the write access to the mygpfs1 file system

The DB2 installer and the instance utilities will not make changes to the GPFS cluster during:

- the creation of a DB2 pureScale instance
- the extension of a DB2 pureScale instance
- the dropping of a DB2 pureScale instance
- or the dropping of a member or cluster caching facility from the instance.

Chapter 15. DB2 Connect considerations with the DB2 pureScale Feature

You cannot install the IBM DB2 pureScale Feature and DB2 Connect products on the same server. If your applications require access to both host and distributed systems, use a separate instance for DB2 Connect functionality and that instance can be a client of the DB2 pureScale server.

The following DB2 Connect server products cannot be installed on a DB2 pureScale server:

- DB2 Connect Enterprise Edition
- DB2 Connect Application Server Edition
- DB2 Connect Unlimited Edition for System i
- DB2 Connect Unlimited Edition for System z

Chapter 16. DB2 client considerations for the DB2 pureScale Feature

Several client features are available with the IBM DB2 pureScale Feature.

For your application to make full use of the DB2 pureScale features, your DB2 server and DB2 client must be at certain release levels:

Server version	Client version	Features available
DB2 Version 9.8, or later	Version 9.7, Fix Pack 1, or later	Transaction-level and connection-level workload balancing Automatic client reroute based on workload Client affinities
DB2 Version 9.8, or later	Version 9.1, Version 9.5, and Version 9.7 (before Fix Pack 1)	Connection-level workload balancing (transaction-level workload balancing is unavailable) Automatic client reroute based on workload
DB2 Version 10.1	Version 9.5, and Version 9.7 (before Fix Pack 1)	Connection-level workload balancing (transaction-level workload balancing is unavailable) Automatic client reroute based on workload
DB2 Version 10.1	Version 9.7, Fix Pack 1, or later	Transaction-level and connection-level workload balancing Automatic client reroute based on workload Client affinities

Client features

Automatic client reroute

Automatic client reroute is an IBM Data Server feature that redirects client applications from a failed server to another server so the applications can continue their work with minimal interruption.

Client affinities

Client affinities provide an ordered list of members to which the client can connect. There is no consideration for the workload of the members, if the first member is unavailable, or if your client is connected to it and it becomes unavailable, the automatic client reroute feature attempts to connect to the next member in the list.

Workload balancing

Automatic workload balancing (WLB) uses member workload information contained in the server list as returned by a DB2 pureScale database server to enable the client to distribute work in a balanced fashion among members.

During COMMIT and ROLLBACK operations, DB2 pureScale database servers restrict clients from using workload balancing if the following conditions apply:

- The connection uses global variables
- An encrypted password is used
- Open With Hold cursors are used

- Declared temporary tables (DGTT) are used
- A TRANSFORM GROUP has been set
- The SESSION AUTHID has been changed
- PL/SQL packages or SQL/PL modules are used
- Cursor variables are used
- Sequence values are used and DB2_ALLOW_WLB_WITH_SEQUENCES is not enabled
- Created temporary tables (CGTT) with PRESERVE ROWS are used

Clients are also restricted from using workload balancing if dynamic SQL is prepared in a package bound with the KEEP_DYNAMIC YES bind option. This restriction does not apply when preparing statements in a stored procedure or user-defined function, or when preparing statements through non-embedded APIs like CLI, JDBC, ODBC, or .NET. If workload balancing is required, you must ensure that the application re-prepares statements across transactions. Once this requirement is satisfied, you must rebind the application with KEEP_DYNAMIC NO.

Applications using the KEEP_DYNAMIC YES bind option in a DB2 pureScale environment behave as follows:

- Re-preparing statements in a new transaction can happen implicitly for applications that are using IBM non-embedded APIs. The existing remote application works as usual.
- For all applications that do not use IBM non-embedded APIs, there is no guarantee that statements will be reprepared in a new transaction. Therefore, the server disallows automatic workload balancing (WLB).

KEEP_DYNAMIC NO, as specified as a bind option or through the ALTER PACKAGE statement, indicates to the server that statements will be reprepared in a new transaction. With this setting the KEEP_DYNAMIC behavior has no bearing on whether WLB will be allowed or disallowed. For applications using IBM embedded APIs, if you can ensure statements will be reprepared in a new transaction, you can employ the KEEP_DYNAMIC NO behavior to allow WLB.

For further information about using these client features, search the DB2 Information Center for information about client high availability connections to DB2 database servers.

Chapter 17. Preparing to install the DB2 pureScale Feature for DB2 Enterprise Server Edition

The following topics provide detailed instructions on how to prepare your environment for the installation of the DB2 pureScale Feature. The DB2 pureScale Feature is supported on specific AIX and Linux operating systems.

Installing DB2 pureScale Feature (AIX)

The topics in this section detail preparing your environment for the installation of the DB2 pureScale Feature in the supported AIX environments.

Installation prerequisites for DB2 pureScale Feature (AIX)

Before you install IBM DB2 pureScale Feature, you must ensure that your system meets the following hardware, firmware, software, and storage requirements. You can use the **db2prereqcheck** command to check the software and firmware prerequisites of a specific DB2 version.

Hardware and firmware prerequisites

The hardware and firmware requirements for supported servers are listed in these tables later in this section:

- Server-specific hardware details:
 - DDR - InfiniBand support
 - QDR - InfiniBand support
 - RDMA over Converged Ethernet (RoCE) support
- Cable information:
 - 10GE cable information (1, 3 and 5 meters)
 - QDR IB cable information (1, 3, 5, 10, 30 meters)
- Switches:
 - Supported InfiniBand network switches
 - IBM validated 10GE switches

Servers in a DB2 pureScale environment must use both an Ethernet network and high-speed communication adapter port.

The communication adapter port can be:

- an InfiniBand (IB) network.
- a 10 Gigabit Ethernet (10GE) network.

To use a 10GE network all 10GE network adapters and switches must be capable of remote direct memory access (RDMA) over Converged Ethernet (RoCE).

Note: Although a single Ethernet adapter is required on a host for the public network in a DB2 pureScale environment, you should set up Ethernet bonding for the network if you have two Ethernet adapters. Ethernet bonding (also known as channel bonding) is a setup where two or more network interfaces are combined. Ethernet bonding provides redundancy and better resilience in the event of Ethernet network adapter failures. Refer to your Ethernet adapter documentation for instructions on configuring Ethernet bonding.

Table 23. Server-specific hardware details for DDR - InfiniBand support

Server	Required Platform Firmware level	InfiniBand network adapter, GX Dual-port 12x Channel Attach - DDR InfiniBand Channel adapter	InfiniBand Channel conversion cables
IBM POWER7 795 (9119-FHB) *	AH720_102 or higher	Feature Code 1816	12x to 4x (Feature Code 1828, 1841, or 1854)
IBM POWER7 780 (9179-MHB) *	AM720_102 or higher	Feature Code 1808	12x to 4x (Feature Code 1828, 1841, or 1854)
IBM POWER7 780 (9179-MHC) *	AM740_042 or higher	Feature Code 1808	12x to 4x (Feature Code 1828, 1841, or 1854)
IBM POWER7 770 (9117-MMB) *	AM720_102 or higher	Feature Code 1808	12x to 4x (Feature Code 1828, 1841, or 1854)
IBM POWER7 770 (9117-MMC) *	AM740_042 or higher	Feature Code 1808	12x to 4x (Feature Code 1828, 1841, or 1854)
IBM POWER7 750 (8233-E8B)	AL730_035 or higher	Feature Code 5609	12x to 4x (Feature Code 1828, 1841, or 1854)
IBM POWER7 740 (8205-E6C)	AL720_102 or higher	Feature Code EJ04	12x to 4x (Feature Code 1828, 1841, or 1854)
IBM POWER7 740 (8205-E6B)	AL720_102 or higher	Feature Code 5615	12x to 4x (Feature Code 1828, 1841, or 1854)
IBM POWER7 730 (8231-E2B)	AL720_102 or higher	Feature Code 5266	4x to 4x cables (Feature Code 3246)
IBM POWER7 720 (8202-E4C)	AL720_102 or higher	Feature Code EJ04	12x to 4x (Feature Code 1828, 1841, or 1854)
IBM POWER7 720 (8202-E4B)	AL720_102 or higher	Feature Code 5615	12x to 4x (Feature Code 1828, 1841, or 1854)
IBM POWER7 710 (8231-E2B)	AL720_102 or higher	Feature Code 5266	4x to 4x cables (Feature Code 3246)
IBM POWER6 [®] 595 (9119-FHA)	EH350_071 or higher	Feature Code 1816	12x to 4x (Feature Code 1828, 1841, or 1854)
IBM POWER6 550 Express (8204-E8A)	EL350_071 or higher	Feature Code 5609	12x to 4x (Feature Code 1828, 1841, or 1854)

Note:

- When acquiring systems, consider the I/O ports available and future workloads for greater flexibility and scalability. The servers marked with an asterisk (*) are designed for enterprise applications. For more information about selecting the

hardware, see “Site and hardware planning” in the IBM System Hardware documentation: <http://publib.boulder.ibm.com/infocenter/powersys/v3r1m5/index.jsp>.

- InfiniBand Channel conversion cables are available in multiple lengths, each with a different product feature code (FC). Some different 12x to 4x InfiniBand Channel conversion cable lengths available are 1.5 m (FC 1828), 3 m (FC 1841), and 10 m (FC 1854). Your data center layout and the relative location of the hardware in the DB2 pureScale environment are factors that must be considered when selecting the cable length.

Table 24. Server-specific hardware details for QDR - InfiniBand support

Server	Required Platform Firmware level	InfiniBand network adapter, GX Dual-port 12x Channel Attach - DDR InfiniBand Channel adapter
IBM POWER7 780/HE (9179-MHC)	AM740_042_042	PCIe2 2-port 4X InfiniBand QDR Adapter (Feature code: 5285)
IBM POWER7 770/MR (9117-MMC)	AM740_042_042	PCIe2 2-port 4X InfiniBand QDR Adapter (Feature code: 5285)
IBM POWER7 740 2S (8205-E6C with optional low-profile slots)	AL740_043_042	PCIe2 2-port 4X InfiniBand QDR Adapter (Feature code: 5285), or PCIe2 Low Profile 2-port 4X InfiniBand QDR Adapter (Feature code: 5283) in the PCIe Newcombe Riser Card (Feature code: 5685), or both
IBM POWER7 740 (8205-E6B) with Newcombe (optional low-profile Gen2 slots)	AL720_102	PCIe2 Low Profile 2-port 4X InfiniBand QDR Adapter (Feature code: 5283) in the PCIe Newcombe Riser Card (Feature code: 5685)
IBM POWER7 710 (8231-E1C)	AL740_043_042	PCIe2 Low Profile 2-port 4X InfiniBand QDR Adapter (Feature code: 5283) in the PCIe Newcombe Riser Card (Feature Code: 5685)
IBM POWER7 720 (8202-E4B)	AL730_066_035	PCIe2 Low Profile 2-port 4X InfiniBand QDR Adapter (Feature code: 5283) in the PCIe Newcombe Riser Card (Feature Code: 5685)
IBM POWER7 720 (8202-E4C)	AL740_043_042	PCIe2 2-port 4X InfiniBand QDR Adapter (Feature code: 5285), or PCIe2 Low Profile 2-port 4X InfiniBand QDR Adapter (Feature code: 5283) in the PCIe Newcombe Riser Card (Feature code: 5685), or both

Table 24. Server-specific hardware details for QDR - InfiniBand support (continued)

Server	Required Platform Firmware level	InfiniBand network adapter, GX Dual-port 12x Channel Attach - DDR InfiniBand Channel adapter
IBM POWER7 730 2S (8231-E2C)	AL740_043_042	PCIe2 Low Profile 2-port 4X InfiniBand QDR Adapter (Feature code: 5283) in the PCIe Newcombe Riser Card (Feature Code: 5685)

Note: QDR IB adapters do not support virtualization. Each LPAR requires a dedicated QDR IB adapter. For example, if a machine is has two LPARs (one for CF and one for member), each of these LPARs must have its own dedicated QDR IB adapter.

Table 25. Server-specific hardware details - RDMA over Converged Ethernet (RoCE) support

Server	Required Platform Firmware level	PCIe Support for RoCE network adapters	10GE Channel conversion cables
IBM POWER7 780/HE (9179-MHC)	AM740_042_042	PCIe2 (Integrated Base Ethernet Adapter and add-on adapters) 10GE card Part Number 00E1493 (FRU) Feature code EC28 (PCIe2 2-Port 10GbE RoCE SFP+ Adapter)	IBM Qualified Copper SFP+ cables (1, 3 and 5 meters)
IBM POWER7 770/MR (9117-MMC)	AM740_042_042	PCIe2 (Integrated Base Ethernet Adapter and add-on adapters) 10GE card Part Number 00E1493 (FRU) Feature code EC28 (PCIe2 2-Port 10GbE RoCE SFP+ Adapter)	IBM Qualified Copper SFP+ cables (1, 3 and 5 meters)

Table 25. Server-specific hardware details - RDMA over Converged Ethernet (RoCE) support (continued)

Server	Required Platform Firmware level	PCIe Support for RoCE network adapters	10GE Channel conversion cables
IBM POWER7 720 1S (8202-E4C with optional low-profile slots)	AL740_043_042	<p>PCIe2 (Integrated Base Ethernet Adapter and add-on adapters)</p> <p>10GE card Part Number 00E1493 (FRU)</p> <p>PCIe2 2-Port 10GbE RoCE SFP+ Adapter (Feature code EC28), or, PCIe2 Low Profile 2-Port 10GbE RoCE SFP+ Adapter (Feature code EC27) in the PCIe Newcombe Riser Card (Feature code 5685)</p>	IBM Qualified Copper SFP+ cables (1, 3 and 5 meters)
IBM POWER7 740 2S (8205-E6C with optional low-profile slots)	AL740_043_042	<p>PCIe2 (Integrated Base Ethernet Adapter and add-on adapters)</p> <p>10GE card Part Number 00E1493 (FRU)</p> <p>PCIe2 2-Port 10GbE RoCE SFP+ Adapter (Feature code EC28), or, PCIe2 Low Profile 2-Port 10GbE RoCE SFP+ Adapter (Feature code EC27) in the PCIe Newcombe Riser Card (Feature code 5685)</p>	IBM Qualified Copper SFP+ cables (1, 3 and 5 meters)
IBM POWER7 710 1S (8231-E1C)	AL740_043_042	<p>PCIe2 (Integrated Base Ethernet Adapter and add-on adapters)</p> <p>10GE card Part Number 00E1493 (FRU)</p> <p>PCIe2 Low Profile 2-Port 10GbE RoCE SFP+ Adapter (Feature code EC27)</p>	IBM Qualified Copper SFP+ cables (1, 3 and 5 meters)

Table 25. Server-specific hardware details - RDMA over Converged Ethernet (RoCE) support (continued)

Server	Required Platform Firmware level	PCIe Support for RoCE network adapters	10GE Channel conversion cables
IBM POWER7 730 2S (8231-E2C)	AL740_043_042	PCIe2 (Integrated Base Ethernet Adapter and add-on adapters) 10GE card Part Number 00E1493 (FRU) PCIe2 Low Profile 2-Port 10GbE RoCE SFP+ Adapter (Feature code EC27)	IBM Qualified Copper SFP+ cables (1, 3 and 5 meters)

Note: RoCE adapters do not support virtualization. Each LPAR requires a dedicated RoCE adapter. For example, if a machine is has two LPARs (one for CF and one for member), each of these LPARs must have its own dedicated RoCE adapter.

Table 26. 10GE cable information (1, 3 and 5 meters)

	1 meter	3 meter	5 meter
IBM Part Number	46K6182	46K6183	46K6184
Feature Code number	EN01	EN02	EN03

Note:

- For 10GE networks you need IBM Qualified Copper 10GE cables that use Small Form-factor Pluggable Plus (SFP+) connections

Table 27. QDR IB cable information (1, 3, 5, 10, 30 meters)

	1 meter (copper)	3 meter (copper)	5 meter (copper)	10 meter (optical)	30 meter (optical)
Feature Code number	3287	3288	3289	3290	3293

Table 28. Supported InfiniBand network switches

InfiniBand switch	Number of ports	Type	Required rack space
IBM 7874-024	24	4x DDR InfiniBand Edge Switch	1U
IBM 7874-040	48	4x DDR InfiniBand Fabric Director Switch	4U
IBM 7874-120	128	4x DDR InfiniBand Fabric Director Switch	7U
IBM 7874-240	288	4x DDR InfiniBand Fabric Director Switch	14U

Table 28. Supported InfiniBand network switches (continued)

InfiniBand switch	Number of ports	Type	Required rack space
IBM 7874-036	36	QDR InfiniBand Switch	1U
IBM 7874-072	72	QDR InfiniBand Switch	5U
IBM 7874-324	324	QDR InfiniBand Switch	14U

Note:

If using two switches in the DB2 pureScale environment, two or more 4x to 4x inter-switch links (ISL) are required. To help with performance and fault tolerance to inter-switch link failures, use half the number of inter-switch link cables as there are total communication adapter ports connected from CFs and members to the switches. For example, in a two switch DB2 pureScale environment where the primary and secondary CF each has four cluster interconnect netnames, and there are four members, use 6 inter-switch links ($6 = (2 \times 4 + 4) / 2$). Choose 4x to 4x InfiniBand ISL cables of appropriate length for your network environment.

Any CX4 to CX4 copper or optical InfiniBand conversion cable can be used for an ISL cable for configurations with DDR IB. For example, any of the Gore & Associates cables identified by the following SKUs can be used: IBN4600-6, IBN4800-5, IBN5800-1, IBN58002, IBN5800-3, and IBN5800-5.

DDR and QDR InfiniBand network switch types cannot be intermixed. The same type of switch, adapter and cables must be used in a cluster. A server using a DDR IB adapter must use a DDR type switch and the corresponding cables. A server using a QDR IB adapter must use a QDR type switch and the corresponding cables.

Table 29. IBM validated 10GE switches

Switch	IBM Validated Switch	Cabling
10GE	Blade Network Technologies® RackSwitch G8124	IBM Qualified Copper SFP+ cables (1, 3 and 5 meters)

Note:

In general, any 10GE switch that supports global pause flow control, as specified by IEEE 802.3x is also supported. However, the exact setup instructions might differ from what is documented in the switch section, which is based on the IBM validated switches. Refer to the switch user manual for details.

Storage hardware requirements

DB2 pureScale Feature supports all storage area network (SAN) and directly attached shared block storage. Configuring DB2 cluster services managed shared storage is recommended for better resiliency. For more information about DB2 cluster services support, see the “Shared storage considerations” topic. The following storage hardware requirements must be met for DB2 pureScale Feature support.

- The following local free disk space on each host:

- 3 GB to extract the installation
 - 3.5 GB for the Installation path
 - 5 GB for the /tmp directory
 - 1 GB for the instance home directory
 - 5 GB for the /var directory.
 - 1 GB for the /(root filesystem)
 - A minimum of three shared file systems are required, each on a separate physical disk. A fourth shared disk is recommended to configure as the DB2 cluster services tiebreaker disk.
- The following shared disk space must be free for each file system:
- Instance shared files: 10 GB²
 - Data: dependent on your specific application needs
 - Logs: dependent on the expectant number of transactions and the applications logging requirements

Note: If the host memory is not enough, you can install the product but cannot start the database instance. Memory requirement is different based on the total number of the database or instances existing on the same host.

Software prerequisites

The AIX operating system version levels, technology levels, and required fixes are listed by supported adapters and switch types:

- DDR IB adapters and switches
- QDR IB adapters and switches
- RoCE adapters and switches

Before running the installation, or the **installFixPack** command, ensure that fixes are applied for your operating system.

Table 30. Specific software requirements for DDR InfiniBand adapters and switches - AIX operating system version and technology levels

AIX version	Technology Level	Minimum Service Pack (SP) level	Required uDAPL level	AIX APAR
AIX 6.1	6	5	6.1.6.15	IV01781 (included in SP6) [*]
AIX 6.1	7	1	6.1.7.0	
AIX 6.1	7	3	6.1.7.1	
AIX 7.1	0	3	7.1.0.15	IV02136 (included in SP4) [*]
AIX 7.1	1	1	7.1.1.0	
AIX 7.1	1	3	7.1.1.2	

2. For better I/O performance, create a separate GPFS file system to hold your database and specify this shared disk on the create database command.

Table 31. Specific software requirements for QDR InfiniBand adapters and switches - AIX operating system version and technology levels

AIX version	Technology Level	Minimum Service Pack (SP) level	Required uDAPL level	AIX APAR
AIX 6.1	7	5	6.1.7.15	IV26383 ¹
AIX 7.1	1	5	7.1.1.15	IV26384 ²

Note:

1. InfiniBand IV26383.epkg.Z ifix is available via FTP from: <ftp://public.dhe.ibm.com/aix/ifixes/iv26383/>.
2. InfiniBand IV26384.epkg.Z ifix is available via FTP from: <ftp://public.dhe.ibm.com/aix/ifixes/iv26384/>.

Table 32. Specific software requirements for RoCE - AIX operating system version and technology levels

AIX version	Technology Level	Minimum Service Pack (SP) level	Required uDAPL level	AIX APAR
AIX 6.1	7	3	6.1.7.2 (updated with IV20240)	IV20240
AIX 7.1	1	3	7.1.1.2 (updated with IV20241)	IV20241

Note:

- Download and install the uDAPL package at the base Technology Level (not the uDAPL packages specific to a fix pack) from the AIX Web Download Pack Programs website, see https://www14.software.ibm.com/webapp/iwm/web/reg/signup.do?source=aixbp&lang=en_US&S_PKG=udapl. After installing the base uDAPL package, apply the appropriate uDAPL fix for the Technology Level from the IBM Support portal, see https://www-304.ibm.com/support/docview.wss?q1=U830315&dc=DB510&rs=1209&uid=isg1fileset664799651&cs=UTF-8&lang=en&loc=en_US.
- If the AIX system is running on a Technology Level with the minimum Service Pack specified in the table, all APARs listed in the row must be installed except for AIX APARs marked with an asterisk (*). The asterisk marked APARs are only required for DB2 pureScale environments with multiple switches. For a system that runs on a Technology Level with a later Service Pack, verify whether the APAR fix is included in the Service Pack level. The first Service Pack the APAR fix is included in is in the table next to the APAR in parentheses. To obtain fixes for the APARs for a system running a Service Pack higher than the minimum required and lower than the Service Pack the fix was first included in, see *IBM Support Fix Central*: <http://www-933.ibm.com/support/fixcentral/>.
- On RoCE networks, after the AIX APAR is applied, the uDAPL level will change.

Required Software:

- OpenSSH level 4.5.0.5302 or higher
- XL C/C++ runtime library level is 9.0.0.12 or higher

Note:

- AIX workload partitions (WPARs) are not supported in a DB2 pureScale environment.
- If you have IBM General Parallel File System (GPFS) already installed, the installation of DB2 pureScale Feature performs the update to the required level automatically. If the existing level is below 3.2.0.0, the installation of DB2 pureScale Feature will not install the required GPFS level automatically, users must perform a manual upgrade of GPFS to the required level. If the existing level is higher than the required level, the installation of DB2 pureScale Feature will skip the GPFS update and use the existing version.
- On AIX systems with existing GPFS installations, you must ensure the GPFS fix g340p11e1 is installed. The fix is found on the DB2 pureScale image in the following directory *<image-directory>/db2/aix/gpfs/efix*.
- If you have IBM Tivoli System Automation for Multiplatforms (Tivoli SA MP) already installed, it must be Tivoli SA MP Version 3.2.2.1. The installation of DB2 pureScale Feature upgrades existing Tivoli SA MP installations to this version level.

Required users for a DB2 pureScale Feature installation

A DB2 pureScale environment requires a user ID for the instance owner, a non-root user ID to use a secure shell (SSH) network protocol between hosts, and another to run fenced user-defined functions or procedures.

The users are required on all servers hosting a cluster caching facility or member. Each user must be configured to have the same user settings and password as the same user on all other servers.

Before you begin

- You must have root user authority to create users and groups.
- If you manage users and groups with security software, additional steps might be required when defining DB2 database users and groups.

About this task

You need these users to create the DB2 pureScale instance:

- One user for the instance owner
- A non-root user ID to use a secure shell (SSH) network protocol between hosts in a DB2 pureScale cluster (this user ID can default to be the instance owner)
- One user for the fenced user

For the instance owner and the fenced user, specify two distinct users that are member of separate groups. You can use existing users and groups, or have them created on your behalf by the installation. If you want the users to be created for you during the installation, the users and groups must not exist on any of the hosts. To use existing users, both users must exist on all the hosts with the same user ID (UID), group ID (GID), and HOME directory before installation.

A DB2 managed GPFS requires a non-root user ID to SSH between hosts in the DB2 pureScale cluster. The user ID must be a user without special privileges. It is used by the root account on the local host to execute commands as root on a remote host in the same cluster. This user ID must be created on all hosts.

The user and group names used in this task are the defaults, and are documented in the following table. You can specify your own user and group names as long as they adhere to your system naming rules and DB2 naming rules.

Table 33. Default users and groups

Required user	User name	Group name
Instance owner	<i>db2sdin1</i>	<i>db2iadm1</i>
Fenced user	<i>db2sdfe1</i>	<i>db2fadm1</i>
Non-root user ID to SSH between hosts	<i>db2ssh1</i>	<i>db2iadm1</i>

If you plan to use the DB2 Setup wizard to install your DB2 database product, the DB2 Setup wizard creates these users for you.

Restrictions

The user names you create must conform to both to the naming rules of your operating system, and to the rules of the DB2 database system.

Procedure

To create these users, perform the following steps:

1. Log on to a host.
2. Create a group for the instance owner and a group for fenced users to run user-defined functions (UDFs) or stored procedures. Before adding the groups on the first host, select a group ID for each group that does not exist on any of the servers.

Operating System	Step
AIX	<p>Run the mkgroup command with the id parameter to create the groups <i>db2iadm1</i> and <i>db2fadm1</i> with the group ID 999 and 998:</p> <pre>mkgroup id=999 db2iadm1 mkgroup id=998 db2fadm1</pre>
Linux	<p>Run groupadd with the -gid option to create the groups <i>db2iadm1</i> and <i>db2fadm1</i> with the group identifiers 999 and 998:</p> <pre>groupadd -gid 999 db2iadm1 groupadd -gid 998 db2fadm1</pre>

3. Create a user that belongs to each group that you created in the previous step. Before adding the users on the first host, choose user ID numbers for both users that do not exist on any of the servers.

Operating System	Step
AIX	<p>Run mkuser to create each user and configure user ID (id), the primary group for the user (pgrp), the complete list of groups the user belongs to (groups), and the home directory of user (home).</p> <pre>mkuser id=1004 pgrp=db2iadm1 groups=db2iadm1 home=/db2home/db2sdin1 core=- 1 data=491519 stack=32767 rss=-1 fsize=-1 db2sdin1 mkuser id=1003 pgrp=db2fadm1 groups=db2fadm1 home=/db2home/db2sdfe1 db2sdfe1</pre>
Linux	<p>Run useradd to create each user, assign the user ID, group, and create the home directory for the user.</p> <pre>useradd -uid 1004 -g db2iadm1 -m -d /db2home/db2sdin1 db2sdin1 useradd -uid 1003 -g db2fadm1 -m -d /db2home/db2sdfe1 db2sdfe1</pre>

This example uses the default instance owner user name, *db2sdin1*, and the default fenced user name, *db2sdfe1*.

- Set an initial password for each user that you created.

```
passwd db2sdin1
passwd db2sdfe1
```

- Log out.
- Log back on with each of the two users you created. You might be prompted to change the user password for each user because it is the first time the user has logged in.
- Log out.
- Create the exact same user and group accounts on each computer that will participate in your database environment.

Installing and setting up OpenSSH

This task describes how to obtain and setup Open Secure Shell (OpenSSH).

Before you begin

Uncomment the entries in the ssh configuration files.

```
File: /etc/ssh/ssh_config
Port 22
Protocol 2
```

```
File: /etc/ssh/sshd_config
PermitRootLogin yes
PasswordAuthentication no
```

About this task

You need to perform these steps on each host you want to participate in the DB2 pureScale instance.

For Linux users, OpenSSH is installed by default on SLES 10 SP3 or later , and RHEL 5.5..

Open Secure Shell (OpenSSH) is an open source version of the SSH protocol suite of network connectivity tools. The tools provide shell functions that are authenticated and encrypted. A shell is a command language interpreter that reads input from a command-line string, stdin, or a file. The steps in this topic will allow you to connect to a remote server through ssh without having to enter a password.

For the DB2 pureScale Feature, you need to set up passwordless SSH access for the root user. The instance owner requires passwordless SSH access, however, the DB2 installation processes set this up if the instance owner does not have it setup. The following steps outline how to set up passwordless SSH access for the root user.

Note: The instance owner must have the rlogin attribute set to the default value of TRUE.

Procedure

1. AIX operating systems only: If OpenSSH is not available on your system, you can obtain it from the latest AIX Expansion Pack and Web Download Pack (<http://www.ibm.com/systems/power/software/aix/expansionpack/index.html>). The OpenSSH fileset includes manual pages with `openssh.man.en_US`. On the Internet, openBSD provides documentation at <http://www.openssh.org/manual.html>.
2. AIX operating systems only: Install OpenSSH. The default setting on AIX for OpenSSH is public key enabled.
3. Setup public key-based authentication. Public key-based authentication enables a single user ID to log in as that same user ID on each host in the instance without being prompted for a password. Public key-based authentication must be enabled for the root ID to use passwordless SSH.

If the user ID has a `~/.ssh` directory, ensure that it does not allow group or other write access. Ensure that the home directory for the user does not allow group or other write access. SSH views this situation as a security exposure and if the directory permissions are not restrictive enough, it does not allow public key-based authentication.

An existing `~/.ssh` directory is not required as the **ssh-keygen** command creates this directory if it does not exist and set up the appropriate access.

From your `~/.ssh` directory, generate a public key/private key pair:

```
$ ssh-keygen -t dsa
```

Whenever prompted for input, hit Enter to accept the default value. (Ensure that no passphrase is entered, or else SSH challenges each authentication attempt, expecting the same passphrase as a response from the user. However, the DB2 product does not allow remote shell utilities to prompt for additional verification.) This action generates two new files in the `~/.ssh` directory, `id_dsa` (the private key) and `id_dsa.pub` (the public key) for DSA encryption.

4. You must generate the public key on each host and append the contents of each public key from each host to a single file called `authorized_keys`. Then copy the `authorized_keys` file to the user's `$HOME/.ssh` directory on each host and run the **chmod 644 authorized_keys** command.

Configuring communication adapter ports and switches (AIX)

The topics in this section detail how to configure communication adapter ports and switches on supported AIX operating systems. The configuration process is dependent on whether you have an InfiniBand network or 10 Gigabit Ethernet (10GE) network.

Configuring on an InfiniBand network (AIX)

The topics in this section detail how to configure single or multiple communication adapter ports and switches on an InfiniBand network.

Configuring switch failover for a DB2 pureScale environment on an InfiniBand network (AIX):

The configuration procedure detailed in this topic is specific to switches in environments with AIX systems and an InfiniBand (IB) network. Switch failover capability is a high availability feature provided by the switch subnet manager (SM) that can be used in multiple switch environments.

Before you begin

1. Set up your network topology according to one the supported network topologies listed in Chapter 12, “Network topology configuration considerations for DB2 pureScale environments,” on page 145.
2. Power on the switch and connect an RJ11 serial cable or Ethernet cable to the switch.

About this task

The procedure details steps for configuring multiple switches to support switch failover, however configuring a single switch includes all steps except the last. Switch failover capability helps the resiliency, or fault tolerance, of a network. Switch failover helps to reduce the detrimental effects of a switch failure by having another switch become the subnet manager if the switch that is the subnet manager fails. Disabling the subnet manager failback setting helps to reduce the effect that the failure of the subnet manager has on network availability. By disabling subnet manager fail back, the secondary subnet manager remains the subnet manager when the original subnet manager rejoins the network after a failure.

Restrictions

Administrative access is required on the switches.

Procedure

1. Connect a console, for example a notebook computer, to the switch. You can use a serial cable to connect to the switch. Alternatively, if you do not have access to a serial cable you can use an Ethernet cable. Follow the instructions to establish a connection for the cabling method you choose:

Cable	Instructions to establish a connection
Serial cable	<ol style="list-style-type: none">1. Connect a console to the switch with a serial cable.2. Open a terminal session from the console to the switch with the following settings:<ul style="list-style-type: none">• 8 data bits• no parity bits• 1 stop bit• 57.6K baud• VT100 emulation• Flow control = XON/XOFF

Cable	Instructions to establish a connection
Ethernet cable	<ol style="list-style-type: none"> 1. Connect a console to the switch with an Ethernet cable. 2. Create a network connection, or modify an existing connection, to use an IP address on the same subnet as the switch. For example, if the IP address of the switch is 192.168.100.10 and the default gateway is 255.255.255.0, configure your console to have the IP address 192.168.100.9 with 255.255.255.0 as the gateway. If you do not know the IP address and gateway of the switch, see the documentation packaged with the switch for information about the default settings. 3. Verify that you can ping the IP of the switch from the console. 4. Open a telnet session to the switch.

2. Configure the default IP and gateway for each switch.
 - a. Logon to the command-line interface of the switch with the admin user ID and password. For information about the default admin ID and password, see the documentation for the switch. For information about switch name and type, see the installation prerequisites for DB2 pureScale Feature topic.
 - b. Set the IP and subnet mask of the switch. Run the **setChassisIpAddr** command with the **-h** parameter to specify the *IP-address* and the **-m** parameter to specify the *subnet-mask*.

```
setChassisIpAddr -h IP-address -m subnet-mask
```

- c. Set the default route for the switch with the **setDefaultRoute** command to use the default gateway IP.

```
setDefaultRoute -h default-gateway
```

3. Reboot the switches so that they use the new configuration.
4. Get the field replaceable unit (Fru) Global Unique Identifier (GUID) for each switch. You can use the web interface for the switch or the command-line interface (CLI):
 - In the web interface, click View Fru and take note the Fru guid field.
 - In the CLI, run the **captureChassis** command or the **fruInfo** command and take note of the FruGuid field.

The field replaceable unit Global Unique Identifier is required to activate the license key for each switch.

5. Activate the subnet manager license keys. You must activate the subnet manager license keys to allow connections to the switches. For information about activating the subnet manager license keys, see the documentation packaged with your switch.
6. Configure the switches so that the selection of the master subnet manager and standby subnet manager is automatic. Use the web interface or the CLI of the switches to start the subnet manager and configure the subnet manager to start when the switch reboots:

- From the switch CLI, run the commands **smControl start** to and **smConfig startAtBoot**.
- Enter the web interface of the switch by entering its IP address into a browser. Click **subnet manager > control > start** to start the subnet manager. Click **subnet manager > configuration > start at boot** to start the subnet manager when the switch reboots.

If the subnet manager is already running, you might encounter an error message reporting that the subnet manager is running. You can ignore this message:

```
smControl start
Starting the SM...
Error trying to control the Subnet manager.
Subnet manager is running. (master)
```

7. Optional: Verify that the subnet manager is running. Run the **smControl** command with the **status** parameter:

```
smControl status
Subnet manager is running. (master)
```

8. Optional: If using multiple switches, you can change the priority on each switch to disable automatic failback of the subnet manager. In most switches, there are usually two priorities:
 - Switch priority - the switch priority determines which switch is selected as the subnet manager. A switch priority of 0 on all the switches results in the switches electing a subnet manager. Always set the switch priority to 0.
 - Elevated priority - the second priority (referred to as the elevated priority) is used to disable automatic failback to the original subnet manager. If this priority is set to 1, and the subnet manager fails and then comes back online, after the failed switch is available, the switch that took over as the subnet manager continues to be the subnet manager. This setting helps reduce unnecessary network delays that are incurred by failing back to the original subnet manager. Set this to 1.

The steps to set the priorities is different for DDR and QDR InfiniBand switches.

- DDR InfiniBand switch

To set the two priorities, enter:

```
smPriority 0 1
```

- QDR InfiniBand switch

To set the two priorities, modify the XML configuration file to set the priority and elevated priority:

- a. Download the file from the GUI. Go to **Config File Admin > Subnet Manager Config File**.
- b. Right click the name beside **Current Config File**, and save the file.
- c. Open the XML file and change <Priority> to 0, and <ElevatedPriority> to 1. For example:

```
<!-- Priority and Elevated Priority control failover for SM, PM and BM. -->
<!-- Priority is used during initial negotiation, high Priority wins. -->
<!-- ElevatedPriority is assumed by winning master. This can prevent -->
<!-- fallback when previous master comes back on line. -->
<Priority>0</Priority> <!-- 0 to 15, higher wins -->
<ElevatedPriority>1</Priority> <!-- 0 to 15, higher wins -->
```

- d. Save the XML file.

- e. Upload the modified XML file back onto the switch. Click **Browse...** beside the "Upload config file:" field. Select the modified file, and click **Upload**.
- f. To have the configuration file take effect, reboot the switch.

Results

The switch, or switches, are now configured for the DB2 pureScale environment.

Example

This example shows the commands entered into the switch CLI for the steps of the procedure on an IBM 7874 switch to configure switch failover. Configure the IP address and subnet mask, then reboot the switch and get the field replaceable unit GUID of the switch:

```
setChassisIpAddr -h IP-address -m subnet-mask
setDefaultRoute -h default-gateway
reboot
fruInfo
```

Activate the switch. To activate switches that use the QLogic firmware (like the IBM 7874 DDR switches), visit the QLogic support website, <http://support.qlogic.com>, and activate the keys for each switch:

1. Click the "License Key Activation" link from the navigation menu. You might receive a prompt for input on how to handle an untrusted security certificate for the QLogic website. You must accept the certificate to activate the license key.
2. Enter the serial number of the switch you want to activate and click **Continue**. The serial number of the switch is in an envelope packaged with the switch. You might be required to enter an email address so that QLogic can send the license key. Provide the email address of the network administrator responsible for the switch, or forward the email to the network administrator.
3. Apply the license key by using the switch CLI or the web interface:
 - On the CLI of the switch, run the **addkey** command.
 - In the web interface, click **License Keys > Key administration > Add key**, enter the license key and click **Apply**.

```
addkey XVARFW-5AKCQS-HDIWS1-EOCTKW-9J3K82-1
showKeys
-----
Key number: 1
Key:       XVARFW-5AKCQS-HDIWS1-EOCTKW-9J3K82-1
Description: Subnet Manager License
Status:    Active
smControl start
smConfig startAtBoot
smControl status
Subnet manager is running. (master)
smPriority 0 1
```

What to do next

Configure the network settings of the hosts, see "Configuring the network settings of hosts in a DB2 pureScale environment on an InfiniBand network (AIX)" on page 178.

Configuring the network settings of hosts in a DB2 pureScale environment on an InfiniBand network (AIX):

As described in the network topology tables and diagrams, configure the communication adapter ports in pairs, so that the devices with the same device ID (for example, ib0) are on the same subnet.

Before you begin

Ensure that you complete the following tasks:

- Read about supported network topologies for DB2 pureScale environments in Chapter 12, “Network topology configuration considerations for DB2 pureScale environments,” on page 145.
- Ensure that your setup conforms to a supported DB2 pureScale environment listed in the “Installation prerequisites for DB2 pureScale Feature (AIX)” on page 161.
- Download the required uDAPL for your TL level as specified in the installation prerequisites.
- Download the uDAPL fileset update.
- Download the InfiniBand fix.
- Assign InfiniBand logical adapter resources to partitions on the Hardware Management Console and start the partitions.

About this task

You must perform these steps on each host, or LPAR, you want to participate in the DB2 pureScale instance. Cluster caching facilities (CFs) and members support multiple communication adapter ports to help DB2 pureScale environments scale and to help with high availability. One communication adapter port for each CF or member is all that is required, though it is recommended to use more adapter ports to increase bandwidth, add redundancy, and allow the use of multiple switches. This topic guides you through the installation and setup of User Direct Access Programming Library (uDAPL) and InfiniBand on AIX hosts and configuring IP addresses.

Restrictions

Administrative access is required on all DB2 member and CF hosts.

Procedure

1. Log in as root.
2. Ensure that all of the following files are present:
 - udapl
 - U825645.bff
 - iz90166.epkg.Z
 - iv01781.epkg.Z (only required for multiple cluster interconnects)
 - iv26383.epkg.Z (only required for QDR IB)
3. If file /etc/dat.conf was previously setup with the desired values, save the existing copy of dat.conf.
4. If you previously installed AIX software from the same directory that now contains the uDAPL files, remove any .toc file that was created from the previous installation. This step ensures that the **smitty** command recognizes all of the files in the specified directory.

```
/home/db2user1/tmp/pkg/> rm *.toc
```


- From the directory that contains the `udapl` package mentioned in step #2 on page 178, run **smitty install**, select “Install and Update Software”, select “Install Software”, then select the input device as `'.'` (the current directory). Then the Install Software screen is shown:

Install Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]	
* INPUT device / directory for software	.	
* SOFTWARE to install	[_all_latest]	+
PREVIEW only? (install operation will NOT occur)	no	+
COMMIT software updates?	yes	+
SAVE replaced files?	no	+
AUTOMATICALLY install requisite software?	yes	+
EXTEND file systems if space needed?	yes	+
OVERWRITE same or newer versions?	no	+
VERIFY install and check file sizes?	no	+
Include corresponding LANGUAGE filesets?	yes	+
DETAILED output?	no	+
Process multiple volumes?	yes	+
ACCEPT new license agreements?	no	+
PREVIEW new LICENSE agreements?	no	+
WPAR Management		
Perform Operation in Global Environment	yes	+
Perform Operation on Detached WPARs	no	+
Detached WPAR Names	[_all_wpars]	+
Remount Installation Device in WPARs	yes	+
Alternate WPAR Installation Device	[]	

Change ACCEPT new license agreements? to yes.

- From the directory that contains the `*.bff` file mentioned in step #2 on page 178, run **smitty update_all** then select the input device as `'.'`. Then the update software screen is shown:

Update Installed Software to Latest Level (Update All)

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]	
* INPUT device / directory for software	.	
* SOFTWARE to update	_update_all	
PREVIEW only? (update operation will NOT occur)	no	+
COMMIT software updates?	yes	+
SAVE replaced files?	no	+
AUTOMATICALLY install requisite software?	yes	+
EXTEND file systems if space needed?	yes	+
VERIFY install and check file sizes?	no	+
DETAILED output?	no	+
Process multiple volumes?	yes	+
ACCEPT new license agreements?	yes	+
PREVIEW new LICENSE agreements?	no	+
WPAR Management		
Perform Operation in Global Environment	yes	+
Perform Operation on Detached WPARs	no	+
Detached WPAR Names	[_all_wpars]	+
Remount Installation Device in WPARs	yes	+
Alternate WPAR Installation Device	[]	

Change ACCEPT new license agreements? to yes.

- Install the InfiniBand interim fixes, `*.epkg.Z` files, by running the **emgr -e** command:

```
emgr -e iz90166.epkg.Z
emgr -e iv01781.epkg.Z
emgr -e iv26383.epkg.Z (only required for QDR IB)
```

8. Verify the installation of the interim fix by running the following command:

```
emgr -lID STATE LABEL INSTALL TIME UPDATED BY ABSTRACT
1 S IZ90166 07/05/11 09:15:04 IB Fixes for UDAPL on 6.1.6.3
```

9. Verify that your system has the correct uDAPL and InfiniBand file sets. Run the following command, shown with sample output, to verify uDAPL is installed correctly:

```
lsipp -l bos.mp64 devices.chrp.IBM.lhca.rte devices.common.IBM.ib.rte udapl.rte
Filesset          Level State      Description
-----
Path: /usr/lib/objrepos
bos.mp64          6.1.6.3 APPLIED   Base Operating System 64-bit
Multiprocessor Runtime
devices.chrp.IBM.lhca.rte 6.1.6.1 APPLIED   Infiniband Logical HCA Runtime
Environment
devices.common.IBM.ib.rte 6.1.6.2 APPLIED   Infiniband Common Runtime
Environment
udapl.rte         6.1.6.1 EFIXLOCKED APPLIED   uDAPL
Path: /etc/objrepos
bos.mp64          6.1.6.3 APPLIED   Base Operating System 64-bit
Multiprocessor Runtime
devices.chrp.IBM.lhca.rte 6.1.6.1 APPLIED   Infiniband Logical HCA Runtime
Environment
devices.common.IBM.ib.rte 6.1.6.2 APPLIED   Infiniband Common Runtime
Environment
udapl.rte         6.1.6.1 EFIXLOCKED APPLIED   uDAPL
-----
```

The command output varies depending on version, technology level, and service pack level.

10. Reboot the system by running the following command:

```
shutdown -Fr
```

11. Configure the InfiniBand subsystem and set IP addresses:

- a. Configure the InfiniBand subsystem in this substep only if InfiniBand was never set up before on the host. Run the **smitty icm** command:

- 1) Select Add an Infiniband Communication Manager
- 2) Key Enter and wait for the command to complete
- 3) Exit by keying Esc+0

For example,

```
Infiniband Communication Manager Device Name      icm
Minimum Request Retries                          [1]
Maximum Request Retries                          [7]
Minimum Response Time (msec)                     [100]
Maximum Response Time (msec)                     [4300]
Maximum Number of HCA's                          [256]
Maximum Number of Users                          [65000]
Maximum Number of Work Requests                  [65000]
Maximum Number of Service ID's                   [1000]
Maximum Number of Connections                    [65000]
Maximum Number of Records Per Request            [64]
Maximum Queued Exception Notifications Per User  [1000]
Number of MAD buffers per HCA                    [64]
```

- b. Run the **smitty inet** command to configure IP addresses.

- 1) Select Change / Show Characteristics
- 2) Select the first IB Interface, "ib0".
- 3) In the Change / Show an IB Interface panel set the values for the fields "Network Interface Name", "INTERNET ADDRESS (dotted decimal)", "Network Mask", "HCA Adapter", and "Adapter's port number". Other fields can be left as the default values. The following CLI terminal screen contains the settings for an example configuration that uses two switches:

Network Interface Name	ib0
INTERNET ADDRESS (dotted decimal)	[10.222.0.1]
Network MASK (hexadecimal or dotted decimal)	[255.255.255.0]
HCA Adapter	[iba0]
Adapter's port number	[1]
Partition Key	[0xFFFF]
MTU	[65532]
Queue Sizes	[4000]
QKey	[0x1E]
Superpacket	off
Interface Specific Network Options ('NULL' will unset the option)	
rfc1323	[1]
tcp_recvspace	[262144]
tcp_sendspace	[262144]
Current STATE	up
Apply change to DATABASE only	no

Note: In the previous example, the *ib0* interface on the *iba0* HCA is in the 10.222.0.0/24 subnet. To enable multiple communication adapter ports on the cluster caching facility (CF) or member, repeat steps “i” through “iii” of part “b” in step 11 for each communication adapter port on each adapter. Each communication adapter port of a host or LPAR must be on a different subnet. Repeat steps 1 - 10 on the secondary CF such that each network interface shares the subnet of the corresponding interface on the primary CF. Repeat steps 1 - 10 on each member. For each successive IB interface, the number should increase. For example, *ib0* will be the first adapter and the subsequent adapter will be *ib1*. Each successive adapter beyond the one listed on port 0, you will need to enter **smitty inet** and select **add an interface -> add an IB interface** instead of change/show interfaces.

12. Reboot the systems by running the following command on each host:

```
shutdown -Fr
```

13. Ensure that the `/etc/dat.conf` file has a format like the example. Ensure the following are configured correctly: the device (`/dev/iba0`), port (1), and InfiniBand adapter name (*ib0*):

- the device is the name of the communication adapter ports
- the port is noted as the adapter's port number
- the Network Interface name is the InfiniBand adapter name.

No modifications are required unless you want to use a different InfiniBand adapter, port, or interface. If the `/etc/dat.conf` file was previously saved after Step 5, verify that the contents are still equivalent. If not, replace the currently `dat.conf` with the saved copy. For example: In the case of a CF or member, the `/etc/dat.conf` would resemble the following example:

```
hca0 u2.0 nonthreadsafe default /usr/lib/libdap1/libdap12.a(shr_64.o) IBM.1.1
"/dev/iba0 1 ib0" " "
hca1 u2.0 nonthreadsafe default /usr/lib/libdap1/libdap12.a(shr_64.o) IBM.1.1
"/dev/iba0 2 ib1" " "
hca2 u2.0 nonthreadsafe default /usr/lib/libdap1/libdap12.a(shr_64.o) IBM.1.1
"/dev/iba1 1 ib2" " "
hca3 u2.0 nonthreadsafe default /usr/lib/libdap1/libdap12.a(shr_64.o) IBM.1.1
"/dev/iba1 2 ib3" " "
```

14. Update the `/etc/hosts` file on each of the hosts so that for each host in the planned DB2 pureScale environment, the file includes all the IP addresses of all the communication adapter ports for all hosts in the planned environment.

For example, in a planned DB2 pureScale environment with multiple communication adapter ports on the CFs and four members, the `/etc/hosts` configuration file might resemble the following file:

```

10.222.0.1      cf1-ib0.example.com cf1-ib0
10.222.1.1      cf1-ib1.example.com cf1-ib1
10.222.2.1      cf1-ib2.example.com cf1-ib2
10.222.3.1      cf1-ib3.example.com cf1-ib3
10.222.0.2      cf2-ib0.example.com cf2-ib0
10.222.1.2      cf2-ib1.example.com cf2-ib1
10.222.2.2      cf2-ib2.example.com cf2-ib2
10.222.3.2      cf2-ib3.example.com cf2-ib3
10.222.0.101    member1-ib0.example.com member1-ib0
10.222.1.101    member1-ib1.example.com member1-ib1
10.222.0.102    member2-ib0.example.com member2-ib0
10.222.1.102    member2-ib1.example.com member2-ib1
10.222.0.103    member3-ib0.example.com member3-ib0
10.222.1.103    member3-ib1.example.com member3-ib1
10.222.0.104    member4-ib0.example.com member4-ib0
10.222.1.104    member4-ib1.example.com member4-ib1

```

Note:

- In a four member environment that uses only one communication adapter port for each CF and member, the file would look similar to the previous example, but contain only the first IP address of each of the CFs in the previous example.
15. Verify the InfiniBand subsystem. Verify the InfiniBand components are in the Available State: For example, the system output of the following command run on a CF verifies that all four IB interfaces and adapters are available:

```

# lsdev -C | grep ib
ib0      Available      IP over Infiniband Network Interface
iba0     Available      InfiniBand host channel adapter
ib1      Available      IP over Infiniband Network Interface
iba0     Available      InfiniBand host channel adapter
ib2      Available      IP over Infiniband Network Interface
iba1     Available      InfiniBand host channel adapter
ib3      Available      IP over Infiniband Network Interface
iba1     Available      InfiniBand host channel adapter
icm      Available      Infiniband Communication Manager

```

Use the **ibstat -v** command to check the InfiniBand state. Verify that the ports are active and the links are up. This check applies only for the port and interface that were previously identified in `/etc/dat.conf` (by default port 1 on `iba0` and the `ib0` interface):

```

-----
IB PORT 1 INFORMATION (iba0)
-----
Global ID Prefix:          fe.80.00.00.00.00.00
Local ID (LID):            000e
Local Mask Control (LMC):  0000
Logical Port State:        Active
Physical Port State:       Active
Physical Port Physical State: Link Up
Physical Port Speed:       5.0G
Physical Port Width:       4X
Maximum Transmission Unit Capacity: 2048
Current Number of Partition Keys: 1
Partition Key List:
  P_Key[0]:                ffff
Current Number of GUID's:  1
Globally Unique ID List:
  GUID[0]:                 00.02.55.00.40.61.40.00

```

Ping from each IB interfaces to every other IB interfaces in the cluster that are in the same IP subnet to make sure that they are reachable. For example,
ping -I <source IP> <destination IP>

What to do next

Configure the storage for the DB2 pureScale environment.

Configuring on a 10GE network (AIX)

The topics in this section detail how to configure single or multiple communication adapter ports and switches on a 10 Gigabit Ethernet (10GE) network.

Configuring switch failover for a DB2 pureScale environment on a 10GE network (AIX):

Switch failover capability is a high availability feature provided by the configure Link Aggregate Control Protocol (LACP) on the switch.

Before you begin

The configuration procedure for switch failover detailed in this topic applies to a 10 Gigabit Ethernet (10GE) network.

Before you begin:

1. Set up your network topology according to one of the supported network topologies listed in Chapter 12, “Network topology configuration considerations for DB2 pureScale environments,” on page 145.
2. Power on the switch and connect an RJ11 serial cable or Ethernet cable to the switch.

About this task

The procedure details steps for configuring two switches to support switch failover. Switch failover capability helps the resiliency, or fault tolerance, of a network. These steps are not applicable for a single switch setup.

To create a DB2 pureScale environment with multiple switches, you must have multiple cluster interconnects on CF servers and configure switch failover on the switches.

Switch failover configuration on a 10GE network requires the support of Link Aggregate Control Protocol (LACP) on the switch.

Restrictions

1. Administrative access is required on the switches.

Procedure

1. Designate two or more ports on each switch to be used as inter-switch links (ISLs).
2. Connect the cables between the two switches.
3. For all switches in the cluster, login through the GUI interface and configure the following. Refer to the switch manual for detailed steps:
 - Disable the Converged Enhance Ethernet (CEE) feature.
 - Ensure Global Pause (IEEE 802.3x) is enabled. For a BNT switch with firmware level 6.5.2 and higher, port flow control must also be enabled for Global Pause.
 - Disable Spanning Tree Protocol (STP).

4. Login to the switch through the GUI interface and configure the following for the ISL ports. Refer to the switch manual for detailed steps:
 - Enable Link Aggregate Control Protocol (LACP).
 - All ISL ports (on both switches) must be setup as active.

Example

What to do next

Configure the network settings of hosts that you plan to include in the DB2 pureScale environment.

Configuring the network settings of hosts in a DB2 pureScale environment on a 10GE network (AIX):

To configure the network settings, you must install required uDAPL software and configure ICM, associate interconnect netnames with pseudo IP addresses, and add required entries to the Direct Access Transport (DAT) configuration file.

Before you begin

Ensure that you complete the following tasks:

- Read about supported network topologies for DB2 pureScale environments in Chapter 12, “Network topology configuration considerations for DB2 pureScale environments,” on page 145.
- Ensure that your setup conforms to a supported DB2 pureScale environment listed in the “Installation prerequisites for DB2 pureScale Feature (AIX)” on page 161.
- Download the required uDAPL for your TL level as specified in the installation prerequisites.
- Download the uDAPL fileset update.

About this task

You must perform these steps on each host, or LPAR, you want to participate in the DB2 pureScale instance. Cluster caching facilities (CFs) and members support multiple communication adapter ports to help DB2 pureScale environments scale and to help with high availability. One communication adapter port for each CF or member is all that is required, though it is recommended to use more adapter ports to increase bandwidth, add redundancy, and allow the use of multiple switches. This topic guides you through the installation and setup of User Direct Access Programming Library (uDAPL) on AIX hosts and configuring IP addresses.

Procedure

1. Log in as root.
2. Ensure that all of the .bff files are present. For example, on AIX 6.1:
 - U848189.bff
 - U848199.bff
 - U849420.bff
 - U849421.bff
 - U849422.bff
 - U849423.bff
3. If file /etc/dat.conf was previously setup with the desired values, save the existing copy of dat.conf.

4. If you previously installed AIX software from the same directory that now contains the uDAPL files, remove any .toc file that was created from the previous installation. This step ensures that the **smitty** command recognizes all of the files in the specified directory.

```
/home/db2user1/tmp/pkg/> rm *.toc
```

5. From the directory that contains the udapl package mentioned in step #2 on page 184, run **smitty install**, select **Install and Update Software**, select **Install Software**, then select the input device as '.' (the current directory). Then the Install Software screen is shown:

Install Software

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]	
* INPUT device / directory for software	.	
* SOFTWARE to install	[_all_latest]	+
PREVIEW only? (install operation will NOT occur)	no	+
COMMIT software updates?	yes	+
SAVE replaced files?	no	+
AUTOMATICALLY install requisite software?	yes	+
EXTEND file systems if space needed?	yes	+
OVERWRITE same or newer versions?	no	+
VERIFY install and check file sizes?	no	+
Include corresponding LANGUAGE filesets?	yes	+
DETAILED output?	no	+
Process multiple volumes?	yes	+
ACCEPT new license agreements?	no	+
PREVIEW new LICENSE agreements?	no	+
WPAR Management		
Perform Operation in Global Environment	yes	+
Perform Operation on Detached WPARs	no	+
Detached WPAR Names	[_all_wpars]	+
Remount Installation Device in WPARs	yes	+
Alternate WPAR Installation Device	[]	

Change **ACCEPT new license agreements?** to yes.

6. From the directory that contains the *.bff file mentioned in step #2 on page 184, run **smitty update_all** then select the input device as '.'. Then the update software screen is shown:

Update Installed Software to Latest Level (Update All)

Type or select values in entry fields.
Press Enter AFTER making all desired changes.

	[Entry Fields]	
* INPUT device / directory for software	.	
* SOFTWARE to update	_update_all	
PREVIEW only? (update operation will NOT occur)	no	+
COMMIT software updates?	yes	+
SAVE replaced files?	no	+
AUTOMATICALLY install requisite software?	yes	+
EXTEND file systems if space needed?	yes	+
VERIFY install and check file sizes?	no	+
DETAILED output?	no	+
Process multiple volumes?	yes	+
ACCEPT new license agreements?	yes	+
PREVIEW new LICENSE agreements?	no	+
WPAR Management		
Perform Operation in Global Environment	yes	+
Perform Operation on Detached WPARs	no	+
Detached WPAR Names	[_all_wpars]	+
Remount Installation Device in WPARs	yes	+
Alternate WPAR Installation Device	[]	

Change **ACCEPT new license agreements?** to yes.

7. Verify that your system has the correct uDAPL and 10GE file sets. Run the following command, shown with sample output, to verify uDAPL is installed correctly:

```
$ lsipp -l bos.mp64 devices.chrp.IBM.lhca.rte devices.common.IBM.ib.rte devices.pciex.b3154a63.rte devices.pciex.b315506714101604.rte udapl.rte
Fileset              Level State      Description
-----
Path: /usr/lib/objrepos
bos.mp64              6.1.7.3 APPLIED    Base Operating System 64-bit
                    Multiprocessor Runtime
devices.chrp.IBM.lhca.rte 6.1.7.3 APPLIED    Infiniband Logical HCA Runtime
                    Environment
devices.common.IBM.ib.rte 6.1.7.4 APPLIED    Infiniband Common Runtime
                    Environment
devices.pciex.b3154a63.rte 6.1.7.4 APPLIED    4X PCI-E DDR Infiniband Device
                    Driver
devices.pciex.b315506714101604.rte 6.1.7.2 APPLIED    Dual Port 10 Gigabit RDMA
                    Converged Ethernet Adapter
                    (RoCE)
                    uDAPL
udapl.rte              6.1.7.2 APPLIED    uDAPL

Path: /etc/objrepos
bos.mp64              6.1.7.3 APPLIED    Base Operating System 64-bit
                    Multiprocessor Runtime
devices.chrp.IBM.lhca.rte 6.1.7.3 APPLIED    Infiniband Logical HCA Runtime
                    Environment
devices.common.IBM.ib.rte 6.1.7.4 APPLIED    Infiniband Common Runtime
                    Environment
devices.pciex.b3154a63.rte 6.1.7.4 APPLIED    4X PCI-E DDR Infiniband Device
                    Driver
udapl.rte              6.1.7.2 APPLIED    uDAPL
```

The command output varies depending on version, technology level, and service pack level.

8. Reboot the system by running the following command:
shutdown -Fr
9. Configure the 10GE subsystem and set IP addresses:
 - a. Configure the 10GE subsystem in this substep only if 10GE was never set up before on the host. Run the **smitty icm** command:

- 1) Select **Add an InfiniBand Communication Manager**
- 2) Key Enter and wait for the command to complete
- 3) Exit by keying Esc+0

For example,

```
Infiniband Communication Manager Device Name      icm
Minimum Request Retries                           [1]
Maximum Request Retries                           [7]
Minimum Response Time (msec)                       [100]
Maximum Response Time (msec)                       [4300]
Maximum Number of HCA's                            [256]
Maximum Number of Users                            [65000]
Maximum Number of Work Requests                    [65000]
Maximum Number of Service ID's                     [1000]
Maximum Number of Connections                      [65000]
Maximum Number of Records Per Request              [64]
Maximum Queued Exception Notifications Per User    [1000]
Number of MAD buffers per HCA                      [64]
```

10. Reboot the systems by running the following command on each host:
shutdown -Fr
11. You must associate each interconnect netname for a member or CF that will be selected during install with an IPv4 pseudo IP address in /etc/hosts or in the DNS database. Each interconnect netname is associated with a RoCE communication adapter port via the Direct Access Transport (DAT) configuration file in the next step. This pseudo IP address is used only for resolving the netname and for uDAPL purposes, it is not pingable. Each pseudo IP address must be unique. Update the /etc/hosts file on each of the hosts so that for each host in the planned DB2 pureScale environment, the file includes all the pseudo IP addresses of interconnect netnames in the planned environment. For example, in a planned DB2 pureScale environment with multiple communication adapter ports on the CFs and four members, the /etc/hosts configuration file might resemble the following file:


```

10.222.1.1      cf1-en1.example.com cf1-en1
10.222.2.1      cf1-en2.example.com cf1-en2
10.222.3.1      cf1-en3.example.com cf1-en3
10.222.4.1      cf1-en4.example.com cf1-en4

10.222.1.2      cf2-en1.example.com cf2-en1
10.222.2.2      cf2-en2.example.com cf2-en2
10.222.3.2      cf2-en3.example.com cf2-en3
10.222.4.2      cf2-en4.example.com cf2-en4

10.222.1.101    member1-en1.example.com member1-en1
10.222.2.101    member1-en2.example.com member1-en2
10.222.1.102    member2-en1.example.com member2-en1
10.222.2.102    member2-en2.example.com member2-en2

10.222.1.103    member3-en1.example.com member3-en1
10.222.2.103    member3-en2.example.com member3-en2
10.222.1.104    member4-en1.example.com member4-en1
10.222.2.104    member4-en2.example.com member4-en2

```

Note: The pseudo IP addresses of each netname for the CF and member must have a different third octet. All pseudo IP address of members must have the same third octet, which is the same as the third octet for the pseudo IP address associated with the first communication adapter port of each of the CFs and members. In the example above the third octet is "1".

All host names in the example above are not associated with regular Ethernet adapters. These host names are set up only for resolving the netnames and for uDAPL purposes. They are not pingable.

In a four member environment that uses only one communication adapter port for each CF and member, the file would look similar to the previous example, but contain only the first pseudo IP address of each of the CFs in the previous example. An example of this is shown below:

```

10.222.1.1      cf1-en1.example.com cf1-en1

10.222.1.2      cf2-en1.example.com cf2-en1

10.222.1.101    member1-en1.example.com member1-en1
10.222.1.102    member2-en1.example.com member2-en1
10.222.1.103    member3-en1.example.com member3-en1
10.222.1.104    member4-en1.example.com member4-en1

```

12. If the Direct Access Transport (DAT) configuration file `/etc/dat.conf` was previously saved, verify that the contents are still equivalent. If the contents are not still equivalent, replace the currently `dat.conf` with the saved copy. If the `dat.conf` file was not previously setup, edit the `dat.conf` file on each host to add a line to associate each interconnect netname with a uDAPL device and a RoCE Adapter port. The following is an example:

```

hca0 u2.0 nonthreadsafe default /usr/lib/libdap1/libdap12.a(shr_64.o) IBM.1.1
"/dev/roce0 1 hostname-en1" " "

```

The name within quotes ("`/dev/roce0 1 hostname-en1`") is the platform-specific string. This string consists of:

- Adapter special file (`/dev/roce0`)
- port number (`1` or `2`)
- The interconnect netname for the member or CF that will run on this host.

The following format is also supported:

```

hca0 u2.0 nonthreadsafe default /usr/lib/libdap1/libdap12.a(shr_64.o) IBM.1.1
"/dev/roce0 1 10.10.11.131" " "

```

Where `10.10.11.131` is the pseudo IP address corresponding to the netname.

Note: If you are receiving a communication error between the member and CF, it is likely that the system attempted to communicate with an adapter interface that is not set up correctly in the Direct Access Transport (DAT) configuration file for the adapter port.

In the case of a CF or member that uses two communication adapters, each communication adapter having 2 ports, the `/etc/dat.conf` would resemble the following example:

```
hca0 u2.0 nonthreadsafe default /usr/lib/libdap1/libdap12.a(shr_64.o)
IBM.1.1 "/dev/roce0 1 cf1-en1" " "
hca1 u2.0 nonthreadsafe default /usr/lib/libdap1/libdap12.a(shr_64.o)
IBM.1.1 "/dev/roce0 2 cf1-en2" " "
hca2 u2.0 nonthreadsafe default /usr/lib/libdap1/libdap12.a(shr_64.o)
IBM.1.1 "/dev/roce1 1 cf1-en3" " "
hca3 u2.0 nonthreadsafe default /usr/lib/libdap1/libdap12.a(shr_64.o)
IBM.1.1 "/dev/roce1 2 cf1-en4" " "
```

13. Verify the 10GE subsystem. Verify the 10GE components are in the Available State: For example, the system output of the following command run on a host, verifies that all devices are available:

```
# lsdev -C | grep -E "Infiniband|PCIE RDMA"
icm          Available      Infiniband Communication Manager
roce0        Available 02-00    PCIE RDMA over Converged Ethernet RoCE Adapter
                                   (b315506714101604)
```

Use the **ibstat -v** command to check the 10GE state. Verify that the ports are active and the links are up. This check applies only for the port and interface that were previously identified in `/etc/dat.conf` (by default port 1 on `roce0`):

```
-----
ETHERNET PORT 1 INFORMATION (roce0)
-----
```

```
Link State: UP
Link Speed: 10G XFI
Link MTU: 9600
Hardware Address: 00:02:c9:4b:97:b8
GIDS (up to 3 GIDs):
GID0 :00:00:00:00:00:00:00:00:00:00:02:c9:4b:97:b8
GID1 :00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00
GID2 :00:00:00:00:00:00:00:00:00:00:00:00:00:00:00:00
```

14. Ensure Global Pause (IEEE 802.3x) is enabled. For details see the "Configuring switch failover" topic.

Configuring PVIDs for a DB2 pureScale instance (AIX)

The shared disk to be used by the DB2 pureScale instance must have the same physical volume identifier (PVID) configured on all hosts.

About this task

The shared disk on all the hosts must have the same PVID configured for the tiebreaker disk and the disks used for the DB2 managed shared file system and must be accessible on all the hosts for instance setup. You need to perform these steps on each host you want to participate in the DB2 pureScale instance for the tiebreaker disk and the disks used for the DB2 managed shared file system.

Procedure

1. Use the **lspv** command to list existing hdisk names (such as device names or physical volumes) and PVIDs. For example:

```
lspv
hdisk0      00c931e42dcf6ce4      rootvg      active
hdisk1      00c931c447f3d4a9      gpfs
```

hdisk2	00c931c447f3d4f2	None
hdisk3	00c931c447f3d517	None
hdisk4	00c931c447f3d53b	None
hdisk5	00c931c447f3d561	None
hdisk6	00c931c447f3d589	None
hdisk7	00c931c447f3d5ac	None
hdisk8	00c931c447f3d5d6	None
hdisk9	none	None

- To assign a PVID to a disk in the storage array, run the following command:

```
chdev -l <disk_name> -a pv=yes
```

For example, for hdisk9, you would run the following command:

```
chdev -l hdisk9 -a pv=yes
```

which would return the following output for the **lspv** command:

lspv			
hdisk0	00c931e42dcf6ce4	rootvg	active
hdisk1	00c931c447f3d4a9	gpfs	
hdisk2	00c931c447f3d4f2	None	
hdisk3	00c931c447f3d517	None	
hdisk4	00c931c447f3d53b	None	
hdisk5	00c931c447f3d561	None	
hdisk6	00c931c447f3d589	None	
hdisk7	00c931c447f3d5ac	None	
hdisk8	00c931c447f3d5d6	None	
hdisk9	00c931e4b4505651	None	

- On each of the other hosts, run the following command to remove the existing disk name:

```
rmdev -dl <disk_name_for_the_same_shared_disk>
```

which would return the following output for the **lspv** command:

lspv			
hdisk0	00c931e42dcf6ce4	rootvg	active
hdisk1	00c931c447f3d4a9	gpfs	
hdisk2	00c931c447f3d4f2	None	
hdisk3	00c931c447f3d517	None	
hdisk4	00c931c447f3d53b	None	
hdisk5	00c931c447f3d561	None	
hdisk6	00c931c447f3d589	None	
hdisk7	00c931c447f3d5ac	None	
hdisk8	00c931c447f3d5d6	None	

- Run the **cfgmgr** command. When the **cfgmgr** command is run, without a PVID assigned, the system will retrieve the PVID from the storage array (set in step 2). Running the **lspv** command will return the following output:

lspv			
hdisk0	00c931e42dcf6ce4	rootvg	active
hdisk1	00c931c447f3d4a9	gpfs	
hdisk2	00c931c447f3d4f2	None	
hdisk3	00c931c447f3d517	None	
hdisk4	00c931c447f3d53b	None	
hdisk5	00c931c447f3d561	None	
hdisk6	00c931c447f3d589	None	
hdisk7	00c931c447f3d5ac	None	
hdisk8	00c931c447f3d5d6	None	
hdisk9	00c931e4b4505651	None	

Preinstallation checklist for DB2 pureScale Feature (AIX)

Perform the following preinstallation steps and verify them on each host before installing the IBM DB2 pureScale Feature.

Before you install

The following steps must be performed on all hosts. The example command output in this topic is from an AIX 6.1 TL6 system. See the installation prerequisites for output values acceptable for other AIX technology levels.

1. DB2 pureScale instances require specific users and groups, including fenced users. You can create the users before starting the DB2 Setup wizard or have the wizard create them for you as you progress through the panels. If you are not creating or modifying instances you can create the required users after completing the installation.
2. Ensure that the required platform firmware level specified in the installation prerequisites is installed. Use the following command to verify that the platform firmware is installed at the correct version level. The output of the command should resemble the sample output:

```
$ lsmlcode -A
sys0!system:EL350_071 (t) EL350_071 (p) EL350_071 (t)
```

The three digits following “EL” in the last column of the command output indicate the platform firmware level. Verify that the platform firmware level meets the required platform firmware installation prerequisite specific to your server.

3. Ensure that the required AIX Version, Technology Level (TL) and Service Pack (SP) are installed according to the software prerequisites. The **oslevel** command displays the operating system level. The following sample output is returned for an AIX 6.1 TL6 system with SP5:

```
$ oslevel -s
6100-06-05-1115
```

4. Ensure that User Direct Access Programming Library (uDAPL) is installed and configured at the level specified for your system in the software prerequisites. The following example illustrates the commands used to verify the uDAPL configuration, run on a system with AIX Version 6.1 TL6 SP5:

```
$ p.IBM.lhca.rte devices.common.IBM.ib.rte udapl.rte
Fileset                                Level State Description
-----
Path: /usr/lib/objrepos
bos.mp64                               6.1.6.15 APPLIED Base Operating System 64-bit
                                         Multiprocessor Runtime

devices.chrp.IBM.lhca.rte              6.1.6.15 APPLIED Infiniband Logical HCA Runtime
                                         Environment

devices.common.IBM.ib.rte              6.1.6.15 APPLIED Infiniband Common Runtime
                                         Environment

udapl.rte                              6.1.6.15 APPLIED uDAPL

Path: /etc/objrepos
bos.mp64                               6.1.6.15 APPLIED Base Operating System 64-bit
                                         Multiprocessor Runtime
                                         Environment

devices.chrp.IBM.lhca.rte              6.1.6.15 APPLIED Infiniband Logical HCA Runtime
                                         Environment

devices.common.IBM.ib.rte              6.1.6.15 APPLIED Infiniband Common Runtime
                                         Environment

udapl.rte                              6.1.6.15 APPLIED uDAPL
```

5. Ensure that OpenSSH is installed and password-less access for the root user is configured on each host. The following command, shown with sample output, is used to verify OpenSSH is installed:

```
$ ls1pp -la "openssh.*"
Fileset                                Level  State      Description
-----
Path: /usr/lib/objrepos
openssh.base.client                    5.4.0.6100  COMMITTED  Open Secure Shell Commands
openssh.base.server                    5.4.0.6100  COMMITTED  Open Secure Shell Server
openssh.license                        4.7.0.5301  COMMITTED  Open Secure Shell License
openssh.man.en_US                      5.4.0.6100  COMMITTED  Open Secure Shell
Documentation - U.S. English
openssh.msg.EN_US                      5.4.0.6100  COMMITTED  Open Secure Shell Messages -
U.S. English (UTF)
openssh.msg.en_US                      5.4.0.6100  COMMITTED  Open Secure Shell Messages -
U.S. English

Path: /etc/objrepos
openssh.base.client                    5.4.0.6100  COMMITTED  Open Secure Shell Commands
openssh.base.server                    5.4.0.6100  COMMITTED  Open Secure Shell Server
```

6. The following command, and sample output, are used to verify the C++ runtime level:

```
$ ls1pp -l x1C.rte
Fileset                                Level  State      Description
-----
Path: /usr/lib/objrepos
x1C.rte                                11.1.0.1  COMMITTED  XL C/C++ Runtime
```

7. Ensure that the shared disks accessed by all hosts have the same physical volume identifier (PVID) configured. Compare these results between each host in the DB2 pureScale instance. The minimum number of shared disks is three. Based on your storage needs, additional disks might be required.

The following sample shows the listing of the PVIDs on a host. The first column identifies the device name and the second column identifies the PVID. In this example, devices *hdisk2*, *hdisk3*, and *hdisk4* are shared disks. To ensure that each shared disk has the same PVID value on each host, compare the output of the **lspv** command on all hosts.

```
$ lspv
hdisk0      00cc14e22575992d      rootvg      active
hdisk1      00cc14e225ae951a      homevg      active
hdisk2      00cc14e22f4acb71      None
hdisk3      00cc14e22f4bf4c1      None
hdisk4      00cc14e22f4c5bfe      None
```

The shared disks must have the same PVID value on each host. However, the shared disk device names can be different on each host. The PVID values in the second column of output must be the same on each host for the shared disks (rows 3, 4, and 5).

8. Ensure that I/O completion ports (IOCPs) are installed and configured. IOCP is not required for the installation of DB2 pureScale Feature, however, it is recommended for performance purposes. The following two commands verify IOCP is installed and enabled. Ensure that the keyword Available displays in the **lsdev** command output.

```
$ ls1pp -l bos.iocp.rte
Fileset                                Level  State      Description
-----
Path: /usr/lib/objrepos
bos.iocp.rte                            6.1.6.0  COMMITTED  I/O Completion Ports API
Path: /etc/objrepos
bos.iocp.rte                            6.1.6.0  COMMITTED  I/O Completion Ports API

$ lsdev -Cc iocp
iocp0 Available  I/O Completion Ports
```

9. If you are updating an instance from the DB2 pureScale Feature in V9.8 Fix Pack 1 to DB2 Version 10.1, before you begin the update, you must manually update the !REQD entry in the host `netmon.cf` file.
10. Optional. If you want to use a specific set of ports for the Fast Communication Manager (FCM), ensure that the ports are free on all hosts. Otherwise, the installer selects a set of unique ports that is available on all the hosts. Port usage is specified in the `/etc/services` file. The FCM requires a port range of the three mandatory ports plus the value provided for the logical members field. This port range can designate up to 130 hosts (128 member + 2 cluster caching facility.) The default FCM start port is 60000 and must be in the range 1024 - 65535. In addition, two ports are required for cluster caching facilities. These two ports are chosen automatically.
11. Confirm that there is an Ethernet and an additional communication adapter port on each server. This additional communication adapter port must be either an InfiniBand or an RDMA capable 10 Gigabit Ethernet (10GE) network adapter. The following command, shown with sample output, lists all available network adapters. Ensure an Ethernet network adapter (`en0`) and an InfiniBand network adapter (`ib0`) or a 10GE network adapter are listed. If different adapter naming conventions are used, check with the network administrator to verify both Ethernet and InfiniBand network adapters are defined.

To verify an InfiniBand network adapter is setup:

```
$ ifconfig -l
en0 ib0 lo0
```

To verify a 10GE network adapter is setup:

```
# lsdev -C | grep -E "Infiniband|PCIE RDMA"
icm          Available          Infiniband Communication Manager
roce0        Available 02-00      PCIE RDMA over Converged Ethernet RoCE Adapter
                                      (b315506714101604)
```

12. As root, validate ssh access between all hosts. From the current host, run the **hostname** command on the current host and on all other hosts in the cluster by using the **ssh** command. The result of the **hostname** command matching the host name identified in the **ssh** command verifies ssh access between the two hosts.

```
ssh <host> hostname
```

For example, on current host *hostpib153*, run the **ssh** command to validate itself and the other three hosts (*hostpib154*, *hostpib155*, and *hostpib156*.) The following commands, and sample output, are used to verify ssh access to the other hosts:

```
root@hostpib153:/> hostname
hostpib153
root@hostpib153:/> ssh hostpib153 hostname
hostpib153
root@hostpib153:/> ssh hostpib154 hostname
hostpib154
root@hostpib153:/> ssh hostpib155 hostname
hostpib155
root@hostpib153:/> ssh hostpib156 hostname
hostpib156
```

13. Optional. For DB2 managed GPFS installations, verify the remote shell and remote file copy settings default to **db2locssh** and **db2scp**. For example:

```
usr/lpp/mmfs/bin/mm1scluster
Remote shell command:    /var/db2/db2ssh/db2locssh
Remote file copy command: /var/db2/db2ssh/db2scp
```

14. If upgrading from DB2 Version 9.8 Fix Pack 2 and earlier, ensure that the .update file, located in <db2 instance shared directory>/sqllib_shared/.update, is synchronized correctly after adding or dropping a member or cluster facility. An example of the file location is /db2sd_20110126085343/db2sdin1/sqllib_shared/.update, where <db2 instance shared directory>=*db2sd_20110126085343*.

To ensure correct synchronization, check that all hosts are listed in the .update file and have the following format: *hostname=install path*. If incorrectly formatted, update the file. For example: *machineA=/opt/IBM/db2/V9.8*, where *hostname=machineA* and *install path=/opt/IBM/db2/V9.8*.

15. As root, ensure that the /tmp directory has at least 5 GB of free space. The following command shows the free space in the /tmp directory.

```
root@hostpib153: /> df -m /tmp
Filesystem      MB blocks      Free %Used    Iused %Iused Mounted on
/dev/hd3         2048.00    1687.23   18%     15927     4% /tmp
```

Using the DB2 Setup wizard

To install the DB2 pureScale Feature, you must know the following items. You can enter your values for each of these steps in the preinstallation cheat sheet section that follows.

- The name of the directory where you want to install the DB2 pureScale Feature.
- The device path for the shared disk partition where the DB2 database and the shared configuration files of the DB2 pureScale instance are stored by default. You should use the DB2 cluster file system.
- If using a DB2 cluster services tiebreaker, the device path for the tiebreaker partition. This partition must be different from the shared disk partition specified earlier.

The DB2 cluster services tiebreaker is used in a recovery scenario when your DB2 pureScale instance is split in half due to a communication failure. DB2 cluster services use the tiebreaker to determine which half of the DB2 pureScale instance will remain online. This partition must be at least 25 MB in size and accessible from all hosts.

- The host name, short host name, or IP address for each of the hosts you want to participate in the DB2 pureScale environment.
- Optionally, you can specify the names of the hosts you want to set up as cluster caching facilities rather than accept the default choices provided by the DB2 Setup wizard. You can manually select which hosts act as cluster caching facilities, or have the DB2 installer assign them.

Preinstallation cheat sheet

Enter the appropriate required item value in the "Your Value" field.

Table 34. Preinstallation cheat sheet

Required Item	Your Value	Example
Instance owner/group name		<i>db2sdin1/db2iadm1</i>
Fenced user/group name		<i>db2sdfe1/db2fadm1</i>
Installation directory name		<i>/opt/IBM/db2/V10.1</i>
Shared file system disk		<i>/dev/hdisk12</i>

Table 34. Preinstallation cheat sheet (continued)

Required Item	Your Value	Example
Hosts to include		<i>db2_host01 - db2_host04.</i>
Netname interconnect for each member and CF		<p>InfiniBand network example: <i>db2_<hostname>_ib0</i></p> <p>10 Gigabit Ethernet network example: <i>db2_<hostname>_en1</i></p> <p>Note: <i>db2_<hostname>_en1</i> does not map to a regular ethernet adapter. It must map to the pseudo IP address for the 10GE communication adapter port.</p>

Note: On InfiniBand only, to find the cluster interconnect netname, run the **netstat -i** command on each host.

Table 35. Preinstallation cheat sheet - optional items

Optional Item	Your Value	Example
Tiebreaker disk		<p>On AIX: <i>/dev/hdisk13</i></p> <p>On Linux: <i>/dev/dm-0</i> or <i>/dev/sdc</i></p>
FCM port range		<i>60000 - 60004</i>
cluster caching facilities port range		<i>56000 - 56001</i>
DB2 communication port		<i>50001</i>
Hosts to set up as cluster caching facilities		<i>db2_host03 and db2host04</i>
On InfiniBand, the cluster interconnect netnames of the cluster caching facilities		<p>Primary: <i>db2_<hostname1>_ib0,</i> <i>db2_<hostname1>_ib1,</i> <i>db2_<hostname1>_ib2,</i> <i>db2_<hostname1>_ib3</i></p> <p>Secondary: <i>db2_<hostname2>_ib0,</i> <i>db2_<hostname2>_ib1,</i> <i>db2_<hostname2>_ib2,</i> <i>db2_<hostname2>_ib3</i></p>

Table 35. Preinstallation cheat sheet - optional items (continued)

Optional Item	Your Value	Example
On 10GE, the cluster interconnect netnames of the cluster caching facilities		Primary: db2_<hostname1>_en1, db2_<hostname1>_en2, db2_<hostname1>_en3, db2_<hostname1>_en4 Secondary: db2_<hostname2>_en1, db2_<hostname2>_en2, db2_<hostname2>_en3, db2_<hostname2>_en4
Hosts to set up as members		db2_host01 and db2host02

What to do next

If you completed all the steps in the preinstallation checklist and filled out the cheat sheet, you can proceed directly to the installation section.

Installing DB2 pureScale Feature (Linux)

The topics in this section detail preparing your environment for the installation of the DB2 pureScale Feature in the supported Linux environments.

Installation prerequisites for DB2 pureScale Feature (Linux)

Before you install IBM DB2 pureScale Feature, you must ensure that your system meets the following network, hardware, firmware, storage, and software requirements. You can use the **db2prereqcheck** command to check the software and firmware prerequisites of a specific DB2 version.

Network prerequisites

Two networks are required, one Ethernet network and one high speed communication network. The high speed communication network must be an InfiniBand (IB) network or a 10 Gigabit Ethernet (10GE) network, a mixture of these two networks is not supported.

Note: Although a single Ethernet adapter is required for a DB2 pureScale Feature environment, you should set up Ethernet bonding for the network if you have two Ethernet adapters. Ethernet bonding (also known as channel bonding) is a setup where two or more network interfaces are combined. Ethernet bonding provides redundancy and better resilience in the event of Ethernet network adapter failures. Refer to your Ethernet adapter documentation for instructions on configuring Ethernet bonding. Bonding high speed communication network is not supported.

Table 36. High speed communication adapter requirements rack mounted servers

Communication adapter type	Switch	IBM Validated Switch	Cabling
InfiniBand (IB)	QDR IB	Mellanox part number MIS5030Q-1SFC	QSFP cables

Table 36. High speed communication adapter requirements rack mounted servers (continued)

Communication adapter type	Switch	IBM Validated Switch	Cabling
10 Gigabit Ethernet (10GE)	10GE	<ol style="list-style-type: none"> 1. Blade Network Technologies RackSwitch G8124 2. Cisco Nexus 5596 Unified Ports Switch 	Small Form-factor Pluggable Plus (SFP+) cables

1. DB2 pureScale environments with Linux systems and InfiniBand communication adapter require FabricIT EFM switch based fabric management software. For communication adapter port support on CF servers, the minimum required fabric manager software image that must be installed on the switch is: image-PPC_M405EX-EFM_1.1.2500.img. The switch might not support a direct upgrade path to the minimum version, in which case multiple upgrades are required. For instructions on upgrading the fabric manager software on a specific Mellanox switch, see the Mellanox website: http://www.mellanox.com/content/pages.php?pg=ib_fabricit_efm_management&menu_section=55. Enabling subnet manager (SM) on the switch is mandatory for InfiniBand networks. To create a DB2 pureScale environment with multiple switches, you must have communication adapter on CF servers and configure switch failover on the switches. To support switch failover, see the Mellanox website for instructions on setting up the subnet manager for a high availability domain.
2. Cable considerations:
 - On InfiniBand networks: The QSFP 4 x 4 QDR cables are used to connect hosts to the switch, and for inter-switch links, too. If using two switches, two or more inter-switch links are required. The maximum number of inter-switch links required can be determined by using half of the total communication adapter ports connected from CFs and members to the switches. For example, in a two switch DB2 pureScale environment where the primary and secondary CF each have four communication adapter ports, and there are four members, the maximum number of inter-switch links required is 6 ($6 = (2 * 4 + 4) / 2$).
 - On a 10GE network, the maximum number of ISLs can be further limited by the number of ports supported by the Link Aggregate Communication Protocol (LACP) which is one of the setup required for switch failover. As this value can differ in different switch vendors, refer to the switch manual for any such limitation. For example, the Blade Network Technologies G8124 24 port switch with Blade OS 6.3.2.0, has a limitation of maximum 8 ports in each LACP trunk between the two switches effectively cap the maximum of ISLs to four (4 ports on each switch).
3. In general, any 10GE switch that supports global pause flow control, as specified by IEEE 802.3x is also supported. However, the exact setup instructions might differ from what is documented in the switch section, which is based on the IBM validated switches. Refer to the switch user manual for details.

Table 37. High speed communication adapter requirements for BladeCenter HS22 servers

Communication adapter type	Switch	Cabling
InfiniBand (IB)	Voltaire 40 Gb InfiniBand Switch ¹ , for example part number 46M6005	QSFP cables ²

Table 37. High speed communication adapter requirements for BladeCenter HS22 servers (continued)

Communication adapter type	Switch	Cabling
10 Gigabit Ethernet (10GE) ³	BNT® Virtual Fabric 10 Gb Switch Module for IBM BladeCenter®, for example part number 46C7191	

1. To create a DB2 pureScale environment with multiple switches, set up communication adapter for the CF hosts.
2. Cable considerations:
 - On InfiniBand networks: The QSFP 4 x 4 QDR cables are used to connect hosts to the switch, and for inter-switch links, too. If using two switches, two or more inter-switch links are required. The maximum number of inter-switch links required can be determined by using half of the total communication adapter ports connected from CFs and members to the switches. For example, in a two switch DB2 pureScale environment where the primary and secondary CF each have four communication adapter ports, and there are four members, the maximum number of inter-switch links required is 6 ($6 = (2 * 4 + 4) / 2$). On a 10GE network, the maximum number of ISLs can be further limited by the number of ports supported by the Link Aggregate Communication Protocol (LACP) which is one of the setup required for switch failover. As this value can differ in different switch vendors, refer to the switch manual for any such limitation. For example, the Blade Network Technologies G8124 24 port switch with Blade OS 6.3.2.0, has a limitation of maximum 8 ports in each LACP trunk between the two switches effectively cap the maximum of ISLs to four (4 ports on each switch).
3. For more information about using DB2 pureScale Feature with application cluster transparency in BladeCenter, see this developerWorks article: <http://www.ibm.com/developerworks/data/library/techarticle/dm-1110purescalebladecenter/>.

Note: If a member exists on the same host as a cluster caching facility (CF), the cluster interconnect netname in db2nodes.cfg for the member and CF must be the same.

Hardware and firmware prerequisites

System x® (x64) servers with either of the specified network adapter types are supported for DB2 pureScale Feature:

Table 38. Supported server configurations

Server	10 Gigabit Ethernet (10GE) adapter	Minimum 10GE network adapter firmware version	InfiniBand (IB) Host Channel Adapter (HCA)	Minimum IB HCA firmware version
System x3650 M3	Mellanox ConnectX-2 EN 10 Gigabit Ethernet Adapters with RoCE	2.9.1000	Mellanox ConnectX-2 with Virtual Protocol Interconnect	2.9.1000
System x3650 M4	MT27500 - Mellanox ConnectX-3 EN Dual-port SFP+ 10GbE Adapter - 00W0053 part number	2.10.2322	MT27500 - Mellanox ConnectX-3 VPI QSFP Dual-Port Infiniband card	2.10.2322

Table 38. Supported server configurations (continued)

Server	10 Gigabit Ethernet (10GE) adapter	Minimum 10GE network adapter firmware version	InfiniBand (IB) Host Channel Adapter (HCA)	Minimum IB HCA firmware version
System x3690 X5	Mellanox ConnectX-2 EN 10 Gigabit Ethernet Adapters with RoCE	2.9.1000	Mellanox ConnectX-2 with Virtual Protocol Interconnect	2.9.1000
System x3850 X5	Mellanox ConnectX-2 EN 10 Gigabit Ethernet Adapters with RoCE	2.9.1000	Mellanox ConnectX-2 with Virtual Protocol Interconnect	2.9.1000
BladeCenter HS22 System x blades	Mellanox 2-port 10 Gb Ethernet Expansion Card with RoCE, for example part number 90Y3570	2.9.1000	2-port 40 Gb InfiniBand Card (CFFh), for example part number 46M6001	2.9.1000
BladeCenter HS23 System x blades	Mellanox 2-port 10 Gb Ethernet Expansion Card (CFFh) with RoCE, part number 90Y3570	2.9.1000	2-port 40 Gb InfiniBand Expansion Card (CFFh) - part number 46M6001	2.9.1000
KVM Virtual Machine	Mellanox ConnectX-2 EN 10 Gb Ethernet Adapters with RoCE	2.9.1200	Not supported	N/A

Note:

1. Install the latest supported firmware for your System x server from <http://www.ibm.com/support/us/en/>.
2. KVM-hosted environments for DB2 pureScale are supported on System x3650, x3690, and x3850 servers.

Storage hardware requirements

DB2 pureScale Feature supports all storage area network (SAN) and directly attached shared block storage. Configuring DB2 cluster services managed shared storage is recommended for better resiliency. For more information about DB2 cluster services support, see the “Shared storage considerations” topic. The following storage hardware requirements must be met for DB2 pureScale Feature support.

- The following local free disk space on each host:
 - 3 GB to extract the installation
 - 3.5 GB for the Installation path
 - 5 GB for the /tmp directory
 - 1 GB for the instance home directory
 - 5 GB for the /var directory.
 - 1 GB for the /(root filesystem)
- A minimum of three shared file systems are required, each on a separate physical disk. A fourth shared disk is recommended to configure as the DB2 cluster services tiebreaker disk.

The following shared disk space must be free for each file system:

- Instance shared files: 10 GB³
- Data: dependent on your specific application needs
- Logs: dependent on the expectant number of transactions and the applications logging requirements

Note: If the host memory is not enough, you can install the product but cannot start the database instance. Memory requirement is different host based on the total number of the database or instances existing on the same host.

Software prerequisites

The libraries and additional packages, listed for each specific Linux distribution in the following table are required on the cluster caching facilities and members. The DB2 pureScale Feature does not support Linux virtual machines. Update hosts with the required software before installing DB2 pureScale Feature or updating to the latest fix pack.

Table 39. Minimum Linux software requirements

Linux distribution	Kernel version level	Required packages	OpenFabrics Enterprise Distribution (OFED) package
Red Hat Enterprise Linux (RHEL) 5.6 ¹	2.6.18-194.26.1.el5	libstdc++ (both 32-bit and 64-bit libraries) glibc++ (both 32-bit and 64-bit libraries) cpp gcc gcc-c++ kernel-headers kernel-devel binutilsOpenSSH sg3_utils ntp-4.2.2p1-15.el5	To install OFED on RHEL 5.6 and higher, run a group installation of "OpenFabrics Enterprise Distribution".

3. For better I/O performance, create a separate GPFS file system to hold your database and specify this shared disk on the create database command.

Table 39. Minimum Linux software requirements (continued)

Linux distribution	Kernel version level	Required packages	OpenFabrics Enterprise Distribution (OFED) package
Red Hat Enterprise Linux (RHEL) 6.1 ⁴	2.6.32-131.0.15.el6	<p>For InfiniBand network type (both 32-bit and 64-bit libraries unless specified) :</p> <p>libibcm dapl (64-bit libraries only) ibsim (64-bit libraries only) ibutils (64-bit libraries only) libibverbs librdmacm libcxgb3 libibmad libibumad libipathverbs (64-bit libraries only) libmlx4 libmthca libnes (64-bit libraries only) libmlx4 rdma (no architecture)</p> <p>For 10GE network type (both 32-bit and 64-bit libraries unless specified) :</p> <p>ibibcm dapl (64-bit libraries only) ibsim (64-bit libraries only) ibutils (64-bit libraries only) libibverbs-rocee librdmacm libcxgb3 libibmad libibumad libipathverbs (64-bit libraries only) libmlx4-rocee libmthca libnes (64-bit libraries only) rdma (no architecture)</p> <p>ntp-4.2.4p8-2.el6.x86_64/ntpdate-4.2.4p8-2.el6.x86_64 libstdc++-4.4.5-6.el6.x86_64 libstdc++-4.4.5-6.el6.i686 glibc-2.12-1.25.el6.x86_64 glibc-2.12-1.25.el6.i686 gcc-c++-4.4.5-6.el6.x86_64 gcc-4.4.5-6.el6.x86_64 kernel-2.6.32-131.0.15.el6.x86_64 kernel-devel-2.6.32-131.0.15.el6.x86_64 kernel-headers-2.6.32-131.0.15.el6.x86_64 kernel-firmware-2.6.32-131.0.15.el6.noarch ntp-4.2.4p8-2.el6.x86_64 ntpdate-4.2.4p8-2.el6.x86_64 sg3_utils-1.28-3.el6.x86_64 sg3_utils-libs-1.28-3.el6.x86_64 binutils-2.20.51.0.2-5.20.el6.x86_64 binutils-devel-2.20.51.0.2-5.20.el6.x86_64 openssh-5.3p1-52.el6.x86_64</p> <p>cpp-4.4.5-6.el6.x86_64</p> <p>ksh-20100621-16.el6.x86_64</p>	<p>For InfiniBand network type, run a group installation of "InfiniBand Support" package.</p> <p>For 10GE network type, subscribe to the Red Hat High Performance Network, then run a group install of "InfiniBand Support" package. This automatically installs the "RHEL server High Performance Networking" package which is mandatory for RDMA over Ethernet support on 10GE network.</p>

Table 39. Minimum Linux software requirements (continued)

Linux distribution	Kernel version level	Required packages	OpenFabrics Enterprise Distribution (OFED) package
SUSE Linux Enterprise Server (SLES) 10 ² Service Pack (SP) 3	2.6.16.60-0.69.1-smp ³	libstdc++ (both 32-bit and 64-bit libraries) glibc++ (both 32-bit and 64-bit libraries) cpp gcc gcc-c++ kernel-source binutils OpenSSH scsi*.rpm ntp-4.2.4p8-1.3.28	For SLES 10 SP3 3, to acquire and install the required OFED packages, see technote #1455818 at http://www.ibm.com/support/docview.wss?uid=swg21455818 For SLES 10 SP4 and later service packs, you must install OFED packages from the maintenance repository with additional packages that OFED depends on. For more information about installing OFED on SLES 10, see “Configuring the network settings of hosts for a DB2 pureScale environment on an InfiniBand network (Linux)” on page 208.
SUSE Linux Enterprise Server (SLES) 11 Service Pack 1	2.6.32.36-0.5 ³	libstdc++ (both 32-bit and 64-bit libraries) glibc++ (both 32-bit and 64-bit libraries) cpp gcc gcc-c++ kernel-default kernel-default-devel kernel-default-base kernel-source kernel-syms binutils OpenSSH sg3_utils ntp-4.2.4p8-1.3.28	For information about installing the OFED package and the packages that it depends on, see installing OFED on SLES 11, see “Configuring the network settings of hosts for a DB2 pureScale environment on an InfiniBand network (Linux)” on page 208.

Table 39. Minimum Linux software requirements (continued)

Linux distribution	Kernel version level	Required packages	OpenFabrics Enterprise Distribution (OFED) package
<p>1. On Red Hat Linux:</p> <ul style="list-style-type: none"> For single communication adapter ports at CFs on InfiniBand network, the minimum support level is RHEL 5.6. For multiple communication adapter ports on InfiniBand network and single or multiple communication adapter port at CFs on 10GE network, the minimum support level is RHEL 6.1. <p>i686 which is 32-bit packages might not get installed by default when installing x86_64 server. Make sure that all the 32-bit dependencies are explicitly installed. For example:</p> <pre>libstdc++-4.4.5-6.el6.i686, pam-1.1.1-8.el6.i686, pam_krb5-2.3.11-6.el6.i686, pam-devel-1.1.1-8.el6.i686, pam_pkcs11-0.6.2-11.1.el6.i686, pam_idap-185-8.el6.i686</pre> <p>Alternatively, run the yum command after creating a source from local DVD or after registering to RHN:</p> <pre>yum install *.i686</pre> <p>2. On SLES 10 Service Pack 4, the minimum supported kernel version level is the default kernel (2.6.16.60-0.85.1-smp).</p> <p>3. On SLES 11 SP1, the default kernel (version 2.6.32.12-0.7-default) must be upgraded to version 2.6.32.36-0.5, which requires that the following kernel packages be installed from the SLES maintenance software repository:</p> <pre>kernel-default-2.6.32.36-0.5.2 kernel-default-devel-2.6.32.36-0.5.2 kernel-default-base-2.6.32.36-0.5.2 kernel-source-2.6.32.36-0.5.2 kernel-syms-2.6.32.36-0.5.2</pre> <p>4. In some installations, if Intel TCO WatchDog Timer Driver modules are loaded by default, they should be blacklisted, so that they do not start automatically or conflict with RSCT. To blacklist the modules, edit the following files:</p> <p>a. To verify if the modules are loaded</p> <pre>lsmod grep -i iTCO_wdt; lsmod grep -i iTCO_vendor_support</pre> <p>b. Edit the configuration files:</p> <ul style="list-style-type: none"> On RHEL 5.x and RHEL 6.1, edit file /etc/modprobe.d/blacklist.conf: <pre># RSCT hatsd blacklist iTCO_wdt blacklist iTCO_vendor_support</pre> <ul style="list-style-type: none"> On SLES, edit file /etc/modprobe.d/blacklist: <pre>add blacklist iTCO_wdt blacklist iTCO_vendor_support</pre> <p>Note: The minimum supported host and guest operating system level for KVM Virtualization is RHEL 6.2. Fibre Channel adapters and 10 GE adapters are required by the virtual machines via PCI Passthrough. For instructions on setting up PCI Passthrough of devices for guest VMs, see the Red Hat website: http://docs.redhat.com/docs/en-US/Red_Hat_Enterprise_Linux/6/html/Virtualization_Host_Configuration_and_Guest_Installation_Guide/chap-Virtualization_Host_Configuration_and_Guest_Installation_Guide-PCI_Assignment.html</p>			

Note:

Note:

- If you have IBM General Parallel File System (GPFS) installed, it must be IBM General Parallel File System (GPFS) 3.5.0.4 with the special fixes required for the DB2 pureScale Feature. If you need to upgrade to IBM General Parallel File System (GPFS) 3.5.0.4 or install the GPFS fixes, the required files are found on the DB2 pureScale installation image in the db2/linuxamd64/gpfs directory.
- If you have IBM Tivoli System Automation for Multiplatforms (Tivoli SA MP) installed, it must be Tivoli SA MP Version 3.2.2.1. This version can be installed by running the **installSAM** command from the db2/linuxamd64/tsamp directory of the DB2 pureScale installation image.
- For the latest information about supported Linux distributions, see <http://www.ibm.com/software/data/db2/linux/validate/>.

Required users for a DB2 pureScale Feature installation (Linux)

Two users and groups are required to operate a DB2 databases environment on Linux operating systems.

Before you begin

- You must have root user authority to create users and groups.
- If you manage users and groups with security software, additional steps might be required when defining DB2 users and groups.

About this task

You need two users to create the DB2 pureScale instance:

- One user for the instance owner
- One user for the fenced user

You should use two different users with two different groups. Each of the two users should have the same UID, GID, group name, and the home directory on all the hosts. Ensure that if any of the users to be used exist on any of the hosts, they much have matching properties. It is not necessary to create these required users before you start the installation. You can create these users as you proceed through the panels of the DB2 Setup wizard or specify them in your response file. If existing users are used, they have to exist on all the hosts and meet the listed requirements.

The user and group names used in the following instructions are the defaults, and are documented in the following table. You can specify your own user and group names as long as they adhere to your system naming rules and DB2 naming rules.

Table 40. Default users and groups

Required user	User name	Group name
Instance owner	db2sdin1	db2iadm1
Fenced user	db2sdfe1	db2fadm1

The user and group names used in the following instructions are documented in the following table. You can specify your own user and group names if they adhere to your system naming rules and DB2 naming rules.

If you are planning to use the DB2 Setup wizard to install your DB2 database product, the DB2 Setup wizard will create these users for you.

Restrictions

The user names you create must conform to both your operating system's naming rules, and those of the DB2 database system.

The same user name that you will create on different hosts must have the same HOME directory. However, the user names must not already exist on any host. If existing user names are used, the user names must exist on all hosts with the same user ID (uid), group ID (gid), group name, and HOME directory.

Procedure

To create these users, perform the following steps:

1. Log on to a host.
2. Create a group for the instance owner (for example, db2iadm1) and a group that will run UDFs or stored procedures (for example, db2fadm1) by entering the following commands:

```
groupadd -g 999 db2iadm1
groupadd -g 998 db2fadm1
```

Ensure that the specific numbers you are using do not currently exist on any of the machines.

3. Create a user that belongs to each group that you created in the previous step using the following commands. The home directory for each user will be the DB2 home directory that you previously created and shared (db2home).

```
useradd -u 1004 -g db2iadm1 -m -d /db2home/db2inst1 db2inst1
useradd -u 1003 -g db2fadm1 -m -d /db2home/db2fenc1 db2fenc1
```

4. Set an initial password for each user that you created by entering the following commands:

```
passwd db2inst1    passwd db2fenc1
```

5. Log out.
6. Log on to the primary computer as each user that you created (db2inst1 and db2fenc1). You might be prompted to change each user's password because this is the first time that these users have logged onto the system.
7. Log out.
8. Create the exact same user and group accounts on each computer that will participate in your database environment.

Installing and setting up OpenSSH

This task describes how to obtain and setup Open Secure Shell (OpenSSH).

Before you begin

Uncomment the entries in the ssh configuration files.

```
File: /etc/ssh/ssh_config
Port 22
Protocol 2
```

```
File: /etc/ssh/sshd_config
PermitRootLogin yes
PasswordAuthentication no
```

About this task

You need to perform these steps on each host you want to participate in the DB2 pureScale instance.

For Linux users, OpenSSH is installed by default on SLES 10 SP3 or later , and RHEL 5.5..

Open Secure Shell (OpenSSH) is an open source version of the SSH protocol suite of network connectivity tools. The tools provide shell functions that are authenticated and encrypted. A shell is a command language interpreter that reads

input from a command-line string, stdin, or a file. The steps in this topic will allow you to connect to a remote server through ssh without having to enter a password.

For the DB2 pureScale Feature, you need to set up passwordless SSH access for the root user. The instance owner requires passwordless SSH access, however, the DB2 installation processes set this up if the instance owner does not have it setup. The following steps outline how to set up passwordless SSH access for the root user.

Note: The instance owner must have the rlogin attribute set to the default value of TRUE.

Procedure

1. AIX operating systems only: If OpenSSH is not available on your system, you can obtain it from the latest AIX Expansion Pack and Web Download Pack (<http://www.ibm.com/systems/power/software/aix/expansionpack/index.html>). The OpenSSH fileset includes manual pages with `openssh.man.en_US`. On the Internet, openBSD provides documentation at <http://www.openssh.org/manual.html>.
2. AIX operating systems only: Install OpenSSH. The default setting on AIX for OpenSSH is public key enabled.
3. Setup public key-based authentication. Public key-based authentication enables a single user ID to log in as that same user ID on each host in the instance without being prompted for a password. Public key-based authentication must be enabled for the root ID to use passwordless SSH.

If the user ID has a `~/.ssh` directory, ensure that it does not allow group or other write access. Ensure that the home directory for the user does not allow group or other write access. SSH views this situation as a security exposure and if the directory permissions are not restrictive enough, it does not allow public key-based authentication.

An existing `~/.ssh` directory is not required as the **ssh-keygen** command creates this directory if it does not exist and set up the appropriate access.

From your `~/.ssh` directory, generate a public key/private key pair:

```
$ ssh-keygen -t dsa
```

Whenever prompted for input, hit Enter to accept the default value. (Ensure that no passphrase is entered, or else SSH challenges each authentication attempt, expecting the same passphrase as a response from the user. However, the DB2 product does not allow remote shell utilities to prompt for additional verification.) This action generates two new files in the `~/.ssh` directory, `id_dsa` (the private key) and `id_dsa.pub` (the public key) for DSA encryption.

4. You must generate the public key on each host and append the contents of each public key from each host to a single file called `authorized_keys`. Then copy the `authorized_keys` file to the user's `$HOME/.ssh` directory on each host and run the **chmod 644 authorized_keys** command.

Configuring communication adapter ports and switches (Linux)

The topics in this section detail how to configure communication adapter ports and switches on supported Linux operating systems. The configuration process is dependent on whether you have an InfiniBand network or 10 Gigabit Ethernet (10GE) network.

Configuring on an InfiniBand network (Linux)

The topics in this section detail how to configure single or multiple communication adapter ports and switches on an InfiniBand network.

Configuring switch failover for a DB2 pureScale environment on an InfiniBand network (Linux):

Switch failover capability is a high availability feature provided by the switch subnet manager (SM).

Before you begin

The configuration procedure for switch failover detailed in this topic applies to both SuSE Linux Enterprise Server (SLES) systems and Red Hat Enterprise Linux Server (RHEL) systems on an InfiniBand (IB) network.

This procedure does not apply to IBM BladeCenter environments. To configure switch failover support for BladeCenter deployments no actions are performed on the switch modules but you must configure communication adapter ports on the primary and secondary cluster caching facilities (CFs).

For DB2 pureScale environments that do not include blade servers, complete the following tasks before you begin:

1. Set up your network topology according to one the supported network topologies listed in Chapter 12, “Network topology configuration considerations for DB2 pureScale environments,” on page 145.
2. Power on the switch and connect an RJ11 serial cable or Ethernet cable to the switch.

About this task

The procedure details steps for configuring multiple switches to support switch failover. Switch failover capability helps the resiliency, or fault tolerance, of a network.

To create a DB2 pureScale environment with multiple switches, you must have communication adapter ports on CF servers and configure switch failover on the switches.

DB2 pureScale environments with Linux systems and InfiniBand cluster interconnect networks require FabricIT EFM switch based fabric management software. For multiple communication adapter ports support on CF servers, the installation prerequisites list the minimum fabric manager image that must be installed on the switch (image-PPC_M405EX-EFM_1.1.2500.img). Depending on the fabric manager version required, a direct upgrade path to the minimum required version might not be supported. In cases where a direct upgrade path is not supported you must upgrade the switch fabric manager software image by installing each subsequent version. For instructions on upgrading fabric manager software on a specific Mellanox switch, see the Mellanox website: http://www.mellanox.com/content/pages.php?pg=ib_fabricit_efm_management&menu_section=55.

For single switch setup on InfiniBand networks, enabling the subnet manager (SM) on the switch is mandatory.

Restrictions

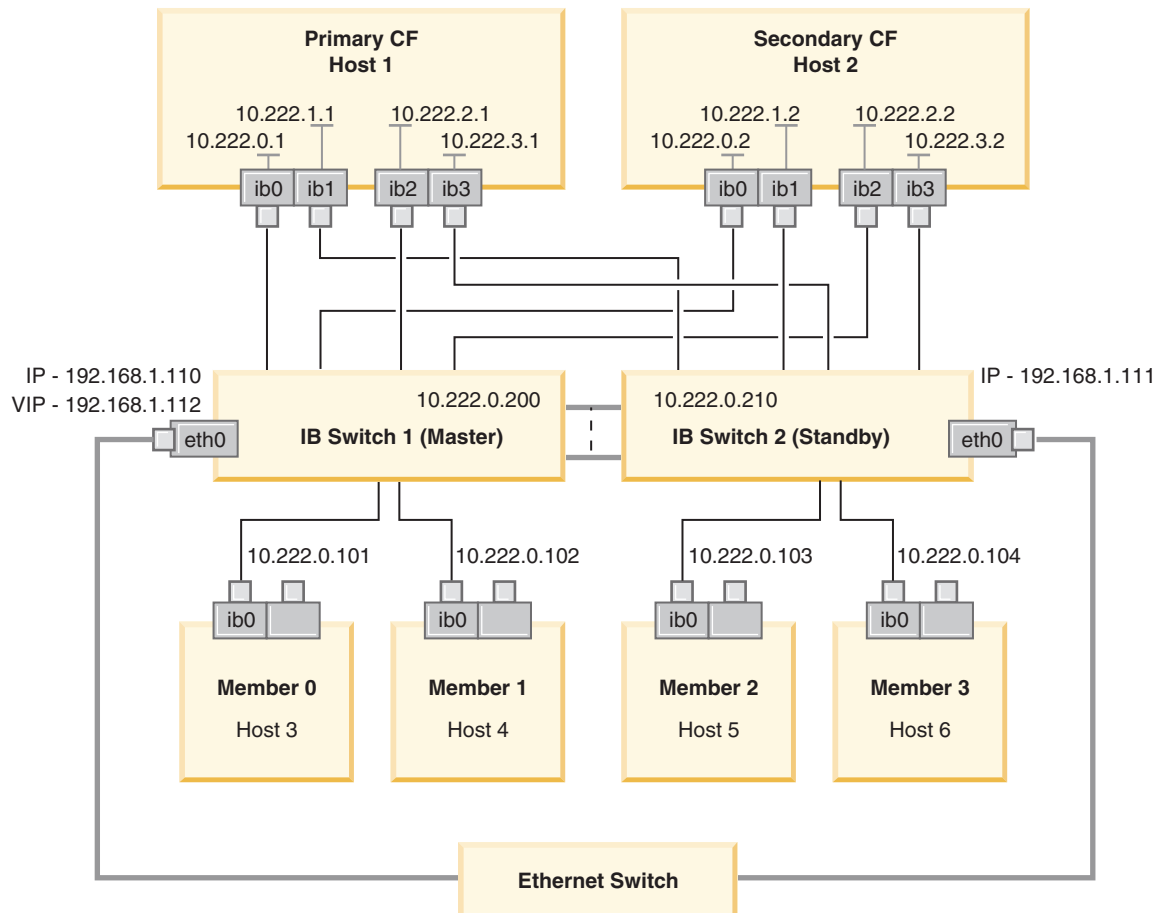
1. Administrative access is required on the switches.

Procedure

1. Upgrade the switch to the required fabric manager software version level listed in the DB2 pureScale Feature installation prerequisites. To obtain Fabric Manager software, see the "Installation prerequisites for DB2 pureScale Feature (Linux)" on page 195.
2. Follow the steps for setting up the Mellanox subnet manager for a high availability domain (Mellanox SM HA), provided in the Mellanox switch user manual.

With the Mellanox SM HA software, the system manager can enter and modify all IB subnet configurations for the subnet managers from a single location. Assign a virtual IP address (VIP) to the management port of the switch to manage the high availability domain. The system manager must configure all the switches in a Mellanox SM HA environment to join the same IB subnet, and assign the subnet a name. After joining the subnet, the subnet managers are synchronized and you must select one as the master subnet manager, and the others become standby subnet managers.

Example



The public Ethernet switch is required to connect the two IB switches to configure the subnet.
Figure 20. Two switch SLES multiple cluster interconnects to CFs

What to do next

Configure the network settings of hosts that you plan to include in the DB2 pureScale environment.

Configuring the network settings of hosts for a DB2 pureScale environment on an InfiniBand network (Linux):

As described in the network topology tables and diagrams, configure the communication adapter ports in pairs, so that the devices with the same device ID (for example, ib0) are on the same subnet.

Before you begin

Ensure that you completed the following tasks:

- Read about supported network topologies for DB2 pureScale environments in Chapter 12, “Network topology configuration considerations for DB2 pureScale environments,” on page 145.
- Ensure that your setup conforms to a supported DB2 pureScale environment listed in the installation prerequisites of Chapter 17, “Preparing to install the DB2 pureScale Feature for DB2 Enterprise Server Edition,” on page 161.

Administrative access is required on all DB2 member and CF hosts.

About this task

To configure the network settings of hosts, install the OpenFabrics Enterprise Distribution (OFED) packages and configure IP addresses on the hosts. Cluster caching facilities (CFs) and members support multiple communication adapter ports to help DB2 pureScale environments scale and to help with high availability. One communication adapter port for each CF or member is all that is required. It is recommended to use more to increase bandwidth, add redundancy, and allow the use of multiple switches.

Note: These steps must be executed on all hosts planned for the future DB2 pureScale environment.

Procedure

1. Log in as root.
2. Configure OpenFabrics Enterprise Distribution (OFED) software.
 - OFED configuration details for SLES systems.
 - For SLES 10 SP3, follow the instructions for installing OFED in technote #1455818 at <http://www.ibm.com/support/docview.wss?uid=swg21455818>
 - For SLES 10 SP4 and later:
 - a. Configure the SLES online updates to include the maintenance repository for your version of SLES.
 - b. Install the following packages from the maintenance repository:
compat-dapl-1.2.19-0.5.1
compat-dapl-32bit-1.2.19-0.5.1
dapl-32bit-2.0.30-0.5.1
dapl-doc-2.0.30-0.5.1
dapl-2.0.30-0.5.1
ibutils-32bit-1.5.4-0.3.3
Ibutils-1.5.4-0.3.3

```

infiniband-diags-1.5.7-0.3.2
libcxgb3-rdmav2-32bit-1.2.5-0.3.1
libcxgb3-rdmav2-1.2.5-0.3.1
libibcm-1.0.5-0.3.1
libibcm-32bit-1.0.5-0.3.1
libibcommon1-1.1.2_20090314-0.1.1
libibcommon1-32bit-1.1.2_20090314-0.1.1
libibmad5-1.3.6-0.3.1
libibmad5-32bit-1.3.6-0.3.1
libibumad3-1.3.6-0.3.1
libibumad3-32bit-1.3.6-0.3.1
libibverbs-1.1.4-0.3.1
libibverbs-32bit-1.1.4-0.3.1
libipathverbs-1.2-0.3.1
libipathverbs-32bit-1.2-0.3.1
libmlx4-rdmav2-1.0-5.21.1
libmlx4-rdmav2-32bit-1.0-5.21.1
libmthca-rdmav2-1.0.5-5.18.1
libmthca-rdmav2-32bit-1.0.5-5.18.1
libnes-rdmav2-1.1.0-0.3.1
librdmacm-1.0.13-0.3.1
librdmacm-32bit-1.0.13-0.3.1
libsdp-32bit-1.1.103-0.3.1
libsdp-1.1.103-0.3.1
mpi-selector-1.0.3-0.3.1
mstflint-1.4-2.25.1
ofed-doc-1.5.2-0.7.1
ofed-kmp-default-1.5.2_2.6.32.29_0.3-0.7.1
ofed-1.5.2-0.7.1
ofed-doc-1.5.2-0.7.1
ofed-kmp-default-1.5.2_2.6.32.29_0.3-0.7.1
opensm-32bit-3.3.7-0.5.1
opensm-3.3.7-0.5.1
ibvexdmtools-0.0.1-75.16.1
qlvnictools-0.0.1-75.16.1
sdpnetstat-1.60-5.22.1
srptools-0.0.4-6.8.2

```

- c. Verify each of the packages that are part of OFED are installed.
- OFED configuration details for RHEL systems.

On RHEL 5.6, run a group installation of the "OpenFabrics Enterprise Distribution" group to install the required InfiniBand packages. Note that on RHEL 5.6, multiple communication adapter ports is not supported. (RHEL 6.1 is required for this) Perform the following as root to install the package:

```
yum groupinstall "OpenFabrics Enterprise Distribution"
```

On RHEL 6.1, run a group installation of the "InfiniBand Support" package to install the required InfiniBand software. The "InfiniBand Support" package is available as a group install. Perform the following as root to install the package:

```
yum groupinstall "InfiniBand Support"
```

Note: For the **yum** command to work, it requires local repositories to be created first from either Red Hat Network (RHN) or from the DVD iso images. After the repository is setup, the **yum** command is aware of the location to find the target packages. Registering with RHN is the recommended mechanisms to access latest kernel updates and fixes. Users are recommended to setup the repository for every RHEL systems. If the repository cannot be setup with RHN, it can also be setup using the iso

images that come with the RHEL DVD media. These procedures are only required on a system if it cannot be registered with RHN.

- a. Copy the file RHEL5.7-20100922.1-Server-x86_64-DVD1.iso from the DVD to a temporary directory on the target system, /tmp/iso

```
# cd /tmp/iso
# ls -rlt
total 3354472
-rw-r--r-- 1 root root 3431618560 Jan 10
20:13 RHEL5.7-20100922.1-Server-x86_64-
DVD1.iso
```

- b. Extract the iso image.

```
mount -o loop /tmp/iso/RHEL5.7-20100922.1
-Server-x86_64-DVD1.iso /mnt/iso/
```

- c. Create a repository.

```
# cd repodata/
# ls -rlt
total 76180
-rw-r--r-- 1 root root 8032315 Jan 17
12:59 primary.xml.gz
-rw-r--r-- 1 root root 51522840 Jan 17
12:59 other.xml.gz
-rw-r--r-- 1 root root 18346363 Jan 17
12:59 filelists.xml.gz
-rw-r--r-- 1 root root 951 Jan 17
12:59 repomd.xml
# cd ..
# cd repodata/
```

- d. Create a repository, by creating a local repository for the iso in /etc/yum.repos.d/my.repo

```
# cat my.repo
[my repo]
name=Redhat LTC
baseurl=file:///mnt/
gpgcheck=0
enabled=1
```

- e. The previous steps complete the creation of the local repository to point to /mnt/iso as the source.
- f. Issue the relevant **yum** command to perform the installation of the required packages.

Sample output for a successful installation:

```
[root@coralxib42 ~]# yum groupinstall 'Infiniband Support'
Loaded plugins: product-id, refresh-packagekit, rhnplugin, subscription-manager
Updating Red Hat repositories.
4/4
Setting up Group Process
Resolving Dependencies
--> Running transaction check
--> Package dap1.x86_64 0:2.0.25-5.2.el6 will be installed
--> Package ibsim.x86_64 0:0.5-4.el6 will be installed
--> Package ibutils.x86_64 0:1.5.4-3.el6 will be installed
--> Processing Dependency: libosmcomp.so.3(OSMCOMP_2.3)(64bit) for package:
ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: libosmvendor.so.3(OSMVENDOR_2.0)(64bit) for package:
ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: libopensm.so.2(OPENSM_1.5)(64bit) for package:
ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: tk for package: ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: libosmcomp.so.3()(64bit) for package:
ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: libosmvendor.so.3()(64bit) for package:
ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: libopensm.so.2()(64bit) for package:
ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: libibdmcom.so.1()(64bit) for package:
```



```

ibutils-1.5.4-3.el6.x86_64
----> Package libcxb3.x86_64 0:1.3.0-1.el6 will be installed
----> Package libibcm.x86_64 0:1.0.5-2.el6 will be installed
----> Package libibmad.x86_64 0:1.3.4-1.el6 will be installed
----> Package libibumad.x86_64 0:1.3.4-1.el6 will be installed
----> Package libibverbs.x86_64 0:1.1.4-4.el6 will be installed
----> Package libibverbs-utils.x86_64 0:1.1.4-4.el6 will be installed
----> Package libipathverbs.x86_64 0:1.2-2.el6 will be installed
----> Package libmlx4.x86_64 0:1.0.1-8.el6 will be installed
----> Package libmthca.x86_64 0:1.0.5-7.el6 will be installed
----> Package libnes.x86_64 0:1.1.1-1.el6 will be installed
----> Package librdmacm.x86_64 0:1.0.10-2.el6 will be installed
----> Package librdmacm-utils.x86_64 0:1.0.10-2.el6 will be installed
----> Package rdma.noarch 0:1.0-9.el6 will be installed
----> Package rds-tools.x86_64 0:2.0.4-3.el6 will be installed
--> Running transaction check
----> Package ibutils-libs.x86_64 0:1.5.4-3.el6 will be installed
----> Package opensm-libs.x86_64 0:3.3.5-1.el6 will be installed
----> Package tk.x86_64 1:8.5.7-5.el6 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package           Arch      Version      Repository      Size
=====
Installing:
dapl x86_64 2.0.25-5.2.el6 rhel-x86_64-server-6 143 k
ibsim x86_64 0.5-4.el6 rhel-x86_64-server-6 55 k
ibutils x86_64 1.5.4-3.el6 rhel-x86_64-server-6 1.0 M
libcxb3 x86_64 1.3.0-1.el6 rhel-x86_64-server-6 16 k
libibcm x86_64 1.0.5-2.el6 rhel-x86_64-server-6 19 k
libibmad x86_64 1.3.4-1.el6 rhel-x86_64-server-6 52 k
libibumad x86_64 1.3.4-1.el6 rhel-x86_64-server-6 55 k
libibverbs x86_64 1.1.4-4.el6 rhel-x86_64-server-6 44 k
libibverbs-utils x86_64 1.1.4-4.el6 rhel-x86_64-server-6 34 k
libipathverbs x86_64 1.2-2.el6 rhel-x86_64-server-6 13 k
libmlx4 x86_64 1.0.1-8.el6 rhel-x86_64-server-6 27 k
libmthca x86_64 1.0.5-7.el6 rhel-x86_64-server-6 33 k
libnes x86_64 1.1.1-1.el6 rhel-x86_64-server-6 15 k
librdmacm x86_64 1.0.10-2.el6 rhel-x86_64-server-6 22 k
librdmacm-utils x86_64 1.0.10-2.el6 rhel-x86_64-server-6 27 k
rdma noarch 1.0-9.el6 rhel-x86_64-server-6 16 k
rds-tools x86_64 2.0.4-3.el6 rhel-x86_64-server-6 55 k
Installing for dependencies:
ibutils-libs x86_64 1.5.4-3.el6 rhel-x86_64-server-6 924 k
opensm-libs x86_64 3.3.5-1.el6 rhel-x86_64-server-6 53 k
tk x86_64 1:8.5.7-5.el6 rhel-x86_64-server-6 1.4 M

Transaction Summary
=====
Install      20 Package(s)

Total download size: 4.0 M
Installed size: 0
Is this ok [y/N]:

```

3. DAT configuration file details for SLES and RHEL systems:

- On SLES, edit the Direct Access Transport (DAT) configuration file, `/etc/dat.conf`, to have a line for each of the communication adapter ports
- On RHEL 5.6 (and higher 5.x releases), the DAT configuration file is located in `/etc/ofed/dat.conf` and it is updated by the group installation of the OFED
- On RHEL 6.1, the DAT configuration file is located in `/etc/rdma/dat.conf` and it is updated by the group installation of the "InfiniBand Support" package

The following example has four 1-port communication adapter ports.

```

cat /etc/dat.conf
ofa-v2-ib0 u2.0 nonthreadsafe default libdaplofa.so.2 dapl.2.0 "ib0 0" ""
ofa-v2-ib1 u2.0 nonthreadsafe default libdaplofa.so.2 dapl.2.0 "ib1 0" ""
ofa-v2-ib2 u2.0 nonthreadsafe default libdaplofa.so.2 dapl.2.0 "ib2 0" ""
ofa-v2-ib3 u2.0 nonthreadsafe default libdaplofa.so.2 dapl.2.0 "ib3 0" ""

```

Note: If you are receiving DAT_INTERNAL_ERR communication errors, it is likely that the system attempted to communicate with an adapter interface that is not set up correctly in the Direct Access Transport (DAT) configuration file for the adapter port.

4. Edit the network configuration files to configure a static IP address for each communication adapter port interface. The following file listings show the network adapter configuration for the CFs, hosts *cf1* and *cf2*, and members, *member1*, *member2*, *member3*, and *member4*. Edit the network configuration files on each host so that the first communication adapter port listed on each host is on the same subnet as the other hosts. If configuring multiple communication adapter ports on the CFs, pair the additional communication adapter ports CFs so that each DEVICE on the secondary CF is on the same subnetwork as the DEVICE on the primary with the same ID.

```
ssh cf1 cat /etc/sysconfig/network-scripts
DEVICE=ib0
BOOTPROTO='static'
IPADDR='10.222.0.1'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'
```

```
ssh cf1 cat /etc/sysconfig/network/ifcfg-ib1
DEVICE=ib1
BOOTPROTO='static'
IPADDR='10.222.1.1'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'
```

```
ssh cf1 cat /etc/sysconfig/network/ifcfg-ib2
DEVICE=ib2
BOOTPROTO='static'
IPADDR='10.222.2.1'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'
```

```
ssh cf1 cat /etc/sysconfig/network/ifcfg-ib3
DEVICE=ib3
BOOTPROTO='static'
IPADDR='10.222.3.1'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'
```

```
ssh cf2 cat /etc/sysconfig/network/ifcfg-ib0
DEVICE=ib0
BOOTPROTO='static'
IPADDR='10.222.0.2'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'
```

```
ssh cf2 cat /etc/sysconfig/network/ifcfg-ib1
DEVICE=ib1
BOOTPROTO='static'
IPADDR='10.222.1.2'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'
```

```
ssh cf2 cat /etc/sysconfig/network/ifcfg-ib2
DEVICE=ib2
BOOTPROTO='static'
IPADDR='10.222.2.2'
```

```

NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'

ssh cf2 cat /etc/sysconfig/network/ifcfg-ib3
DEVICE=ib3
BOOTPROTO='static'
IPADDR='10.222.3.2'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'

ssh member1 cat /etc/sysconfig/network/ifcfg-ib0
DEVICE=ib0
BOOTPROTO='static'
IPADDR='10.222.0.101'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'

ssh member2 cat /etc/sysconfig/network/ifcfg-ib0
DEVICE=ib0
BOOTPROTO='static'
IPADDR='10.222.0.102'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'

ssh member3 cat /etc/sysconfig/network/ifcfg-ib0
DEVICE=ib0
BOOTPROTO='static'
IPADDR='10.222.0.103'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'

ssh member4 cat /etc/sysconfig/network/ifcfg-ib0
DEVICE=ib0
BOOTPROTO='static'
IPADDR='10.222.0.104'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'

```

Note:

- For simplicity, the IP addresses used in the previous example use the 255.255.255.0 subnetwork mask (NETMASK) so that the subnetwork third and forth IP segments can match the numbers of the interface devices and hostname. This subnetwork mask results in the IP addresses for CFs formatted like 10.222.*interface-id-device-number.CF-hostname-suffix* and members IP addresses like 10.222.*interface-id-device-number.10member-hostname-suffix*.
 - The first communication adapter port on each CF host is on the same subnet as the members.
 - Each communication adapter port on a CF or member is on a distinct subnet.
 - Communication adapter ports with the same interface *DEVICE* name on the primary and secondary CFs share the same subnet.
5. If configuring multiple communication adapter ports on members, use the same IP subnet for each adapter interface device on the second host as was used for adapter interface with the same device ID on the other hosts so that matching devices are on the same IP subnets.

```
cat /etc/sysconfig/network/ifcfg-ib0
DEVICE=ib0
BOOTPROTO='static'
IPADDR='10.1.1.161'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'
```

```
cat /etc/sysconfig/network/ifcfg-ib1
DEVICE=ib1
BOOTPROTO='static'
IPADDR='10.1.2.161'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'
```

All members must be on an IP subnet used by the CF adapter interfaces. The resulting IP subnets are:

- The 10.1.1 subnet has the ib0 device from all members and all CFs.
 - The 10.1.2 subnet has the ib1 device from all members and all CFs.
6. For BladeCenter deployments only, enable the subnet manager service (Open SM) on all hosts in the DB2 pureScale environment. To enable the subnet manager service run the following commands on each host to start the service and have it start after a reboot:

```
chkconfig opensmd on
service opensmd start
```

7. Update the /etc/hosts file on each of the hosts so that for each host in the planned DB2 pureScale environment, the file includes all the IP addresses of all the communication adapter ports for all hosts in the planned environment.

For example, in a planned DB2 pureScale environment with multiple communication adapter ports on the CFs with four members, the /etc/hosts configuration file might resemble the following file:

```
10.222.0.1      cf1-ib0.example.com cf1-ib0
10.222.1.1      cf1-ib1.example.com cf1-ib1
10.222.2.1      cf1-ib2.example.com cf1-ib2
10.222.3.1      cf1-ib3.example.com cf1-ib3
10.222.0.2      cf2-ib0.example.com cf2-ib0
10.222.1.2      cf2-ib1.example.com cf2-ib1
10.222.2.2      cf2-ib2.example.com cf2-ib2
10.222.3.2      cf2-ib3.example.com cf2-ib3
10.222.0.101    member1-ib0.example.com member1-ib0
10.222.1.101    member1-ib1.example.com member1-ib1
10.222.0.102    member2-ib0.example.com member2-ib0
10.222.1.102    member2-ib1.example.com member2-ib1
10.222.0.103    member3-ib0.example.com member3-ib0
10.222.1.103    member3-ib1.example.com member3-ib1
10.222.0.104    member4-ib0.example.com member4-ib0
10.222.1.104    member4-ib1.example.com member4-ib1
```

Note:

- In a four member environment that uses a communication adapter port for each CF and member, the file would look similar to the previous example, but contain only the first IP address of each of the CFs in the previous example.
8. Restart the service for the InfiniBand subsystem.

```
service openibd restart
```

On RHEL 6.1:

```
service rdma restart
```

9. Verify the InfiniBand subsystem.

- a. Verify that the ports are active and the links are up. Use the **ibstat -v** command or the **ibstatus** command to list the state of the adapters. This check applies to the ports and interfaces that were previously identified in `/etc/dat.conf`.

```
ibstatus
Infiniband device 'mlx4_0' port 1 status:
  default gid: fe80:0000:0000:0000:0002:c903:0007:eafb
  base lid:    0x2
  sm lid:      0x1
  state:       4: ACTIVE
  phys state:  5: LinkUp
  rate:        20 Gb/sec (4X DDR)

Infiniband device 'mlx4_0' port 2 status:
  default gid: fe80:0000:0000:0000:0002:c903:0007:eafb
  base lid:    0x3
  sm lid:      0x1
  state:       4: ACTIVE
  phys state:  5: LinkUp
  rate:        20 Gb/sec (4X DDR)
```

Note: Port 1 of the example output the **ibstatus** command on Linux corresponds to port 0 in the `dat.conf` file:

```
ofa-v2-ib0 u2.0 nonthreadsafe default libdaplofa.so.2 dapl.2.0 "ib0 0" ""
```

Verify that the state field value is ACTIVE and the phys state field reports that the link is up (LinkUp).

- b. Ensure the destination IP is resolvable. For example, enter the following:

```
# ip -resolve neigh
coralxib44-ib3 dev ib3 lladdr
80:00:00:49:fe:80:00:00:00:00:00:00:00:02:c9:03:00:0e:9d:5e REACHABLE
coralxib42.torolab.ibm.com dev bond0 lladdr 00:1a:64:c9:d1:e8 REACHABLE
coralxib42-ib0 dev ib0 lladdr
80:00:00:48:fe:80:00:00:00:00:00:00:00:02:c9:03:00:07:ea:5f REACHABLE
coralxib44-ib0 dev ib0 lladdr
80:00:00:48:fe:80:00:00:00:00:00:00:00:02:c9:03:00:07:eb:13 REACHABLE
9.26.120.1 dev bond0 lladdr 00:00:0c:07:ac:01 REACHABLE
coralxib43.torolab.ibm.com dev bond0 lladdr 00:1a:64:c9:cc:d4 REACHABLE
coralxib44-ib2 dev ib2 lladdr
80:00:00:48:fe:80:00:00:00:00:00:00:00:02:c9:03:00:0e:9d:5d REACHABLE
coralxib44.torolab.ibm.com dev bond0 lladdr 00:1a:64:c9:d5:24 REACHABLE
coralxib44-ib1 dev ib1 lladdr
80:00:00:49:fe:80:00:00:00:00:00:00:00:02:c9:03:00:07:eb:14 REACHABLE
coralxib43-ib0 dev ib0 lladdr
80:14:00:48:fe:80:00:00:00:00:00:00:00:02:c9:03:00:07:ea:07 REACHABLE

# arp -an
? (10.1.4.144) at 80:00:00:49:fe:80:00:00:00 [infiniband] on ib3
? (9.26.120.241) at 00:1a:64:c9:d1:e8 [ether] on bond0
? (10.1.1.142) at 80:00:00:48:fe:80:00:00:00 [infiniband] on ib0
? (10.1.1.144) at 80:00:00:48:fe:80:00:00:00 [infiniband] on ib0
? (9.26.120.1) at 00:00:0c:07:ac:01 [ether] on bond0
? (9.26.120.103) at 00:1a:64:c9:cc:d4 [ether] on bond0
? (10.1.2.144) at 80:00:00:48:fe:80:00:00:00 [infiniband] on ib2
? (9.26.120.104) at 00:1a:64:c9:d5:24 [ether] on bond0
? (10.1.3.144) at 80:00:00:49:fe:80:00:00:00 [infiniband] on ib1
? (10.1.1.143) at 80:14:00:48:fe:80:00:00:00 [infiniband] on ib0
```

What to do next

Modify the kernel parameters of hosts that you plan to include in the DB2 pureScale environment.

Configuring on a 10GE network (Linux)

The topics in this section detail how to configure single or multiple communication adapter ports and switches on a 10 Gigabit Ethernet (10GE) network.

Setting up the IP interfaces on the switch on a 10GE network (Linux):

On a 10 Gigabit Ethernet (10GE) network, you must manually set up the IP interfaces on the switch.

Before you begin

Ensure that you completed the following tasks:

- The switch to be configured must be turned on.
- You must be able to administrate the switch through a console. For example, a host that is connected to the switch's management port.
- You must have administrative access to the switch.
- The switch must be setup with an IP for its management port and it must be running.

About this task

The number of IP addresses to be setup on a switch is the same as the number of distinct IP subnets directly connected to it from the CF or member server.

Procedure

To set up the IP interfaces

1. Logon to the command-line interface of the switch with the admin user ID and password.
2. Create the interface and setup its IP address using these guidelines:
 - a. Create the same number of IP interfaces on a switch as the number of distinct IP subnets connected from a CF or member host. (For example, in figure 1 below, each switch has two distinct IP subnets connected from the CF host.)
 - b. Each IP interface must be assigned an IP address that is on one of the distinct IP subnets from the CF host.
 - c. No IP subnet can be used more than once on a switch.
3. For a single switch cluster, the setup is now completed. If the cluster in figure 1 only has one switch, after completing steps 1 to 2, the switch will have the following IP addresses assigned to their IP interfaces :
192.168.1.2, 192.168.2.2, 192.168.3.2 and 192.168.4.2.

The final step is to verify the newly created IP interfaces are pingable from each other and each adapter in members and CFs can ping the external pingable IP address(es) specified for that adapter in their **netmon.cf**.

4. For a two-switch cluster, perform the following additional steps :
 - a. Repeat steps 1 and 2 on the other switch. At the end of this step, there should be equal number of IP interfaces on each switch with each IP

address on a different IP subnet. The total number of distinct IP subnets on both switches is the same as the maximum of the total number of distinct IP subnets on a CF or member host.

- b. Create an additional IP interface on the switch where its newly created IP interfaces do not have the same IP subnet as the member's adapter. Assign this new IP interface an IP address that is on the same IP subnet as the member. After this step, one switch should have one extra IP interface than the other. In addition, both switches will have exactly one IP interface where its IP address is on the same IP subnet as all the members.

Using the figure 1 below as an example, after completing steps 1 to 4a, each switch will have the following IP addresses assigned to their IP interfaces :

Switch 1 - 192.168.1.2 and 192.168.2.2

Switch 2 - 192.168.3.2 and 192.168.4.2

Since all members are on the 192.168.1.0 IP subnet, switch 2 will be chosen using step 4's instruction to create the additional IP interface since switch 2 currently does not have an IP address that is on the 192.168.1.0 IP subnet.

After step 4b, each switch will have the following IP addresses assigned :

Switch 1 - 192.168.1.2 and 192.168.2.2

Switch 2 - 192.168.1.5, 192.168.3.2 and 192.168.4.2

Note: The IP interfaces setup procedure is now completed once the 192.168.1.0 IP subnets exist on both switches. The final step is to verify the newly created IP interfaces are pingable from each other and each adapter in members and CFs can ping the external pingable IP addresses specified for that adapter in their **netmon.cf**.

Example

Below is the step-by-step instructions on how to create the IP interfaces on a cluster using two BNT switches :

1. Log in with the admin user ID.
2. Create the interface and setup its IP address and IP subnet mask.
3. If you were creating two interfaces on the switch:
 - a. To create interface 1, setup the IP address and IP subnet mask. For example:

```
RS G8124(config)#interface ip 1
RS G8124(config-ip-if)#ip address 192.168.1.2
RS G8124(config-ip-if)#ip netmask 255.255.255.0
RS G8124(config-ip-if)#exit
```
 - b. To create interface 2 using the same procedure. For example:

```
RS G8124(config)#interface ip 2
RS G8124(config-ip-if)#ip address 192.168.2.2
RS G8124(config-ip-if)#ip netmask 255.255.255.0
RS G8124(config-ip-if)#exit
```
4. For the change to take effect, enter the following command:

```
RS G8124(config)#copy running-config startup-config
```
5. To create interfaces for different IP subnets, repeat steps 2 to 4 on the second switch.
6. To validate the IP addresses, ping the address from the CF and member hosts that connect to the same switch.




Figure 21. Two CFs and four members connect to two switches.

Configuring switch failover for a DB2 pureScale environment on a 10GE network (Linux):

Switch failover capability is a high availability feature provided by the configure Link Aggregate Control Protocol (LACP) on the switch.

Before you begin

The configuration procedure for switch failover detailed in this topic applies to both SuSE Linux Enterprise Server (SLES) systems and Red Hat Enterprise Linux Server (RHEL) systems on a 10 Gigabit Ethernet (10GE) network.

This procedure does not apply to IBM BladeCenter environments. To configure switch failover support for BladeCenter deployments no actions are performed on the switch modules but you must configure multiple cluster interconnects on the primary and secondary cluster caching facilities (CFs) and members.

For DB2 pureScale environments that do not include blade servers, complete the following tasks before you begin:

1. Set up your network topology according to one of the supported network topologies listed in Chapter 12, “Network topology configuration considerations for DB2 pureScale environments,” on page 145.
2. Power on the switch and connect an RJ11 serial cable or Ethernet cable to the switch.

Administrative access is required on the switches.

About this task

The procedure details steps for configuring multiple switches to support switch failover. Switch failover capability helps the resiliency, or fault tolerance, of a network. These steps are not applicable for a single switch setup.

To create a DB2 pureScale environment with multiple switches, you must have multiple cluster interconnects on CF and member servers and configure switch failover on the switches.

Switch failover configuration on a 10GE network requires the support of Link Aggregate Control Protocol (LACP) on the switch.

Procedure

1. Designate two or more ports on each switch to be used as inter-switch links (ISLs).
2. Connect the cables between the two switches.
3. Login to the switch through the GUI interface and configure the following. Refer to the switch manual for detailed steps:
 - All ISL ports must have Link Aggregate Control Protocol (LACP) enabled
 - All ISL ports (on both switches) must be setup as active
 - Disable Spanning Tree Protocol (STP)

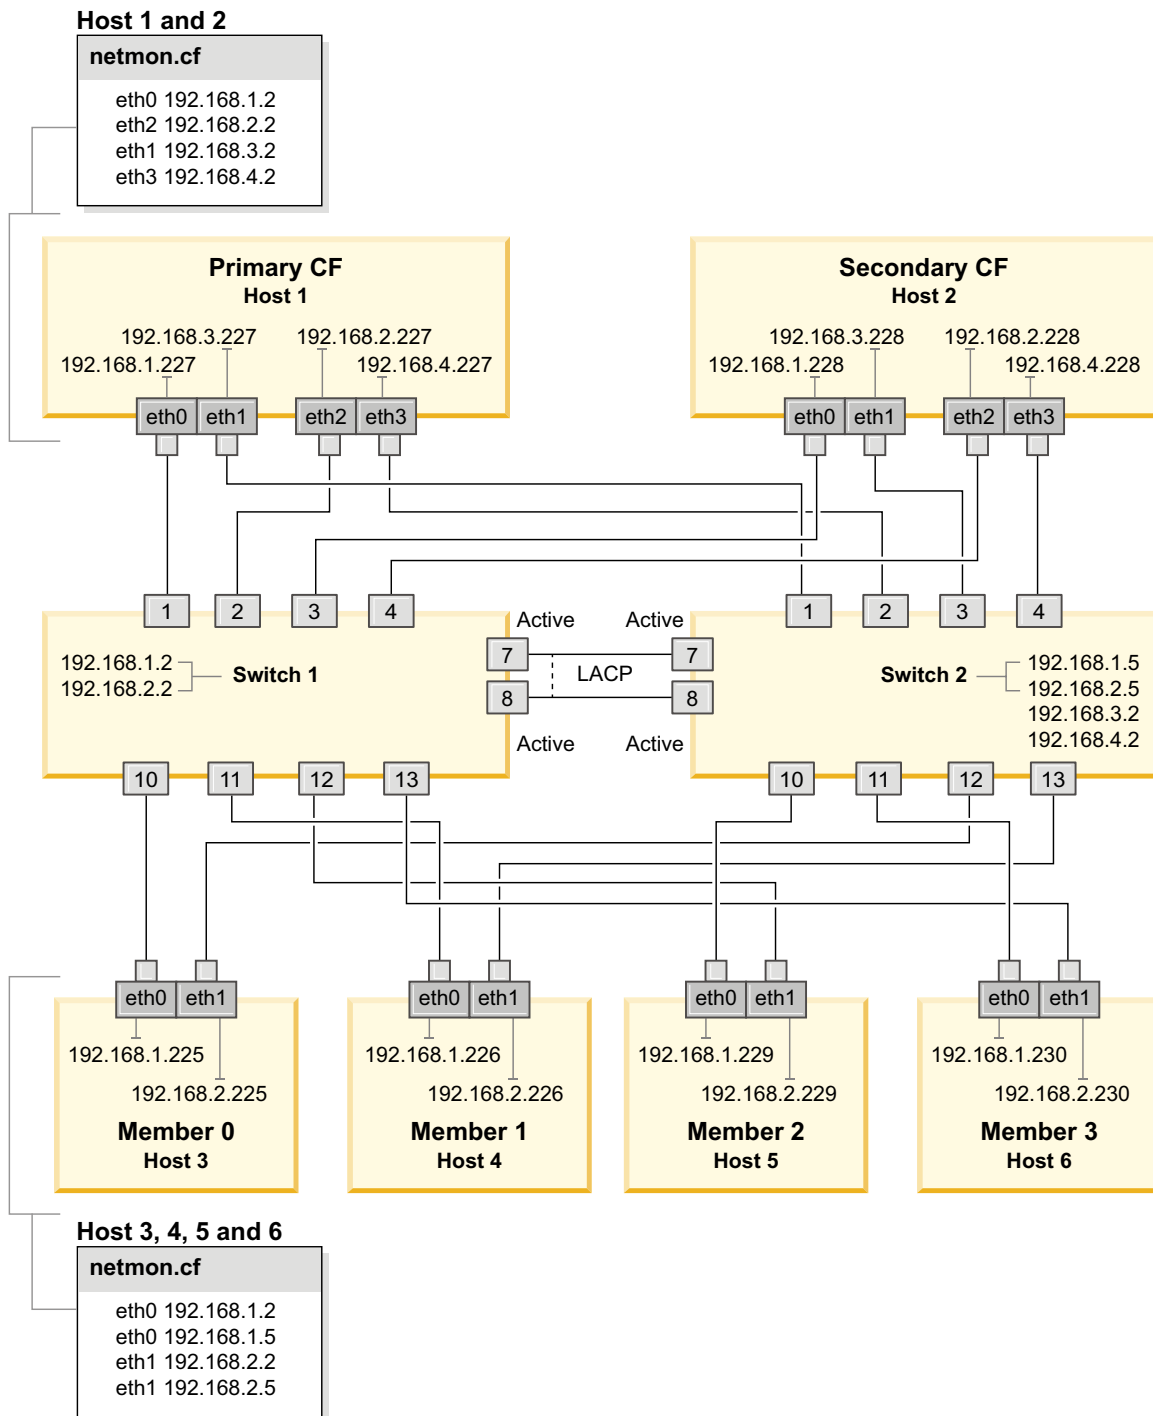


Figure 22. Two CFs and four members connect to two switches.

What to do next

Configure the network settings of hosts that you plan to include in the DB2 pureScale environment.

Configuring the network settings of hosts for a DB2 pureScale environment on a 10GE network (Linux):

As described in the network topology tables and diagrams, configure the communication adapter ports in pairs, so that the devices with the same device ID (for example, ib0) are on the same subnet.

Before you begin

Ensure that you completed the following tasks:

- Read about supported network topologies for DB2 pureScale environments in Chapter 12, “Network topology configuration considerations for DB2 pureScale environments,” on page 145.
- Ensure that your setup conforms to a supported DB2 pureScale environment listed in the installation prerequisites of Chapter 17, “Preparing to install the DB2 pureScale Feature for DB2 Enterprise Server Edition,” on page 161.

Administrative access is required on all DB2 member and CF hosts.

About this task

To configure the network settings of hosts, install the OpenFabrics Enterprise Distribution (OFED) packages on SuSE Linux or the High Performance Networking package on Red Hat Linux, and configure IP addresses on the hosts. Cluster caching facilities (CFs) and members support multiple communication adapter ports to help DB2 pureScale environments scale and to help with high availability. One communication adapter port for each CF or member is all that is required. It is recommended to use more to increase bandwidth, add redundancy, and allow the use of multiple switches.

Note: These steps must be executed on all hosts planned for the future DB2 pureScale environment.

Procedure

1. Log in as root.
2. Configure the appropriate software to support RDMA over the desired network.

For details regarding supported Linux environments, see <https://www.ibm.com/developerworks/wikis/display/im/DB2+and+DB2+Connect+10.1+for+Linux+-+Supported+Environments>.

- OFED configuration details for SLES systems.
 - For SLES 10 SP3, follow the instructions for installing OFED in technote #1455818 at <http://www.ibm.com/support/docview.wss?uid=swg21455818>
 - For SLES 10 SP4 and later :
 - a. Configure the SLES online updates to include the maintenance repository for your version of SLES.
 - b. Install the following packages from the maintenance repository. For example:

```
compat-dapl-1.2.19-0.5.1
compat-dapl-32bit-1.2.19-0.5.1
dapl-32bit-2.0.30-0.5.1
dapl-doc-2.0.30-0.5.1
dapl-2.0.30-0.5.1
ibutils-32bit-1.5.4-0.3.3
Ibutils-1.5.4-0.3.3
infiniband-diags-1.5.7-0.3.2
libcxgb3-rdmav2-32bit-1.2.5-0.3.1
```

```

libcxgb3-rdmav2-1.2.5-0.3.1
libibcm-1.0.5-0.3.1
libibcm-32bit-1.0.5-0.3.1
libibcommon1-1.1.2_20090314-0.1.1
libibcommon1-32bit-1.1.2_20090314
-0.1.1
libibmad5-1.3.6-0.3.1
libibmad5-32bit-1.3.6-0.3.1
libibumad3-1.3.6-0.3.1
libibumad3-32bit-1.3.6-0.3.1
libibverbs-1.1.4-0.3.1
libibverbs-32bit-1.1.4-0.3.1
libipathverbs-1.2-0.3.1
libipathverbs-32bit-1.2-0.3.1
libmlx4-rdmav2-1.0-5.21.1
libmlx4-rdmav2-32bit-1.0-5.21.1
libmthca-rdmav2-1.0.5-5.18.1
libmthca-rdmav2-32bit-1.0.5-5.18.1
libnes-rdmav2-1.1.0-0.3.1
librdmacm-1.0.13-0.3.1
librdmacm-32bit-1.0.13-0.3.1
libsdp-32bit-1.1.103-0.3.1
libsdp-1.1.103-0.3.1
mpi-selector-1.0.3-0.3.1
mstflint-1.4-2.25.1
ofed-doc-1.5.2-0.7.1
ofed-kmp-default-1.5.2_2.6.32.29_0.3
-0.7.1
ofed-1.5.2-0.7.1
ofed-doc-1.5.2-0.7.1
ofed-kmp-default-1.5.2_2.6.32.29_0.3
-0.7.1
opensm-32bit-3.3.7-0.5.1
opensm-3.3.7-0.5.1
ibvexdmttools-0.0.1-75.16.1
qlvnictools-0.0.1-75.16.1
sdpnstat-1.60-5.22.1
srptools-0.0.4-6.8.2

```

c. Verify each of the packages that are part of OFED are installed.

- OFED configuration details for RHEL systems.

On RHEL 5.6, run a group installation of the "OpenFabrics Enterprise Distribution" group to install the required InfiniBand packages. On RHEL 5.6, multiple communication adapter ports in CF is not supported. (RHEL 6.1 is required for this)

On RHEL 6.1:

- Subscribe to RHEL Server High Performance Networking (v. 6 for x86_64) channel. This requires a fee. If already subscribed, there is no need to re-subscribe.
- Run a group installation of the "InfiniBand Support" package to install the required software. This automatically triggers the installation of the RHEL server High Performance Networking package (HPN). A correct installation has the libibverbs and libmlx4 that are packaged with "InfiniBand Support" package replaced with libibverbs-rocee and libibmlx4-rocee.

Sample output:

```

[root@coralm234 ~]# yum groupinstall 'Infiniband Support'
Loaded plugins: product-id, refresh-packagekit, rhnplugin, subscription-manager
Updating Red Hat repositories.
rhel-x86_64-server-hpn-6          9/9
rhel-x86_64-server-hpn-6-debuginfo 4/4
Setting up Group Process
Package libibverbs-utils is obsoleted by libibverbs-rocee-utils, trying
to install libibverbs-rocee-utils-1.1.4-4.el6.x86_64 instead

```

```

Package libibverbs is obsoleted by libibverbs-rocee, trying to install
libibverbs-rocee-1.1.4-4.el6.x86_64 instead
Package libmlx4 is obsoleted by libmlx4-rocee, trying to install
libmlx4-rocee-1.0.1-8.el6.x86_64 instead

Resolving Dependencies
--> Running transaction check
--> Package dapl.x86_64 0:2.0.25-5.2.el6 will be installed
--> Package ibsim.x86_64 0:0.5-4.el6 will be installed
--> Package ibutils.x86_64 0:1.5.4-3.el6 will be installed
--> Processing Dependency: libosmcomp.so.3(OSMCOMP_2.3)(64bit) for package:
ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: libosmvendor.so.3(OSMVENDOR_2.0)(64bit) for package:
ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: libopensm.so.2(OPENSM_1.5)(64bit) for package:
ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: tk for package: ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: libosmcomp.so.3()(64bit) for package:
ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: libosmvendor.so.3()(64bit) for package:
ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: libopensm.so.2()(64bit) for package:
ibutils-1.5.4-3.el6.x86_64
--> Processing Dependency: libibdmcom.so.1()(64bit) for package:
ibutils-1.5.4-3.el6.x86_64
--> Package libcxgb3.x86_64 0:1.3.0-1.el6 will be installed
--> Package libibcm.x86_64 0:1.0.5-2.el6 will be installed
--> Package libibmad.x86_64 0:1.3.4-1.el6 will be installed
--> Package libibumad.x86_64 0:1.3.4-1.el6 will be installed
--> Package libibverbs-rocee.x86_64 0:1.1.4-4.el6 will be installed
--> Package libibverbs-rocee-utils.x86_64 0:1.1.4-4.el6 will be installed
--> Package libipathverbs.x86_64 0:1.2-2.el6 will be installed
--> Package libmlx4-rocee.x86_64 0:1.0.1-8.el6 will be installed
--> Package libmthca.x86_64 0:1.0.5-7.el6 will be installed
--> Package libnes.x86_64 0:1.1.1-1.el6 will be installed
--> Package librdmacm.x86_64 0:1.0.10-2.el6 will be installed
--> Package librdmacm-utils.x86_64 0:1.0.10-2.el6 will be installed
--> Package rdma.noarch 0:1.0-9.el6 will be installed
--> Package rds-tools.x86_64 0:2.0.4-3.el6 will be installed
--> Running transaction check
--> Package ibutils-libs.x86_64 0:1.5.4-3.el6 will be installed
--> Package opensm-libs.x86_64 0:3.3.5-1.el6 will be installed
--> Package tk.x86_64 1:8.5.7-5.el6 will be installed
--> Finished Dependency Resolution

Dependencies Resolved

=====
Package Arch Version Repository Size
=====
Installing:
dapl x86_64 2.0.25-5.2.el6 rhel-x86_64-server-6 143 k
ibsim x86_64 0.5-4.el6 rhel-x86_64-server-6 55 k
ibutils x86_64 1.5.4-3.el6 rhel-x86_64-server-6 1.0 M
libcxgb3 x86_64 1.3.0-1.el6 rhel-x86_64-server-6 16 k
libibcm x86_64 1.0.5-2.el6 rhel-x86_64-server-6 19 k
libibmad x86_64 1.3.4-1.el6 rhel-x86_64-server-6 52 k
libibumad x86_64 1.3.4-1.el6 rhel-x86_64-server-6 55 k
libibverbs-rocee x86_64 1.1.4-4.el6 rhel-x86_64-server-hpn-6 44 k
libibverbs-rocee-utils x86_64 1.1.4-4.el6 rhel-x86_64-server-hpn-6 34 k
libipathverbs x86_64 1.2-2.el6 rhel-x86_64-server-6 13 k
libmlx4-rocee x86_64 1.0.1-8.el6 rhel-x86_64-server-hpn-6 27 k
libmthca x86_64 1.0.5-7.el6 rhel-x86_64-server-6 33 k
libnes x86_64 1.1.1-1.el6 rhel-x86_64-server-6 15 k
librdmacm x86_64 1.0.10-2.el6 rhel-x86_64-server-6 22 k
librdmacm-utils x86_64 1.0.10-2.el6 rhel-x86_64-server-6 27 k
rdma noarch 1.0-9.el6 rhel-x86_64-server-6 16 k
rds-tools x86_64 2.0.4-3.el6 rhel-x86_64-server-6 55 k
Installing for dependencies:
ibutils-libs x86_64 1.5.4-3.el6 rhel-x86_64-server-6 924 k
opensm-libs x86_64 3.3.5-1.el6 rhel-x86_64-server-6 53 k
tk x86_64 1:8.5.7-5.el6 rhel-x86_64-server-6 1.4 M

Transaction Summary
=====
Install 20 Package(s)

Total download size: 4.0 M
Installed size: 0
Is this ok [y/N]:

To verify the RoCE packages installed, enter:
yum groupinfo "High Performance Networking"

```

3. On RHEL 6.1 only, several 32-bit packages required by RSCT must be installed manually as they are no longer installed automatically on RHEL 6.1. The list and the command to install them are shown below. Run these commands as root:

```
yum install libibcm.i686
yum install libibverbs-rocee.i686
yum install librdmacm.i686
yum install libcxgb3.i686
yum install libibmad.i686
yum install libibumad.i686
yum install libmlx4-rocee.i686
yum install libmthca.i686
```

4. Edit the Direct Access Transport (DAT) configuration file to have a line for each of the communication adapter ports. On SLES, the DAT configuration file is located in `/etc/dat.conf`. On RHEL 5.6, the DAT configuration file is located in `/etc/ofed/dat.conf`. On RHEL 6.1, it is located in `/etc/rdma/dat.conf`. This file is updated by the group installation of the packages in previous step. The following is an example of the configuration file on a CF host or member that uses four communication adapter ports:

```
ofa-v2-roe0 u2.0 nonthreadsafe default libdaplofa.so.2 dapl.2.0 "eth0 0" ""
ofa-v2-roe1 u2.0 nonthreadsafe default libdaplofa.so.2 dapl.2.0 "eth1 0" ""
ofa-v2-roe2 u2.0 nonthreadsafe default libdaplofa.so.2 dapl.2.0 "eth2 0" ""
ofa-v2-roe3 u2.0 nonthreadsafe default libdaplofa.so.2 dapl.2.0 "eth3 0" ""
```

Note: If you are receiving `DAT_INTERNAL_ERR` communication errors, it is likely that the system attempted to communicate with an adapter interface that is not set up correctly in the Direct Access Transport (DAT) configuration file for the adapter port.

5. Optional: Install the "infiniband-diags" package for diagnostics utilities (such as `ibstat`, and `ibstatus`). To install the "infiniband-diags" package, run the following command:

```
yum install infiniband-diags
```

The RDMA service must be restarted before running the diagnostics tools. Run the following commands as root:

```
chkconfig rdma on
service rdma restart
```

6. Verify the Global Pause (IEEE 802.3x) flow control is enabled in the 10GE adapter driver. For example, to verify in the Mellanox Connect X-2 10GE adapter driver, the priority bit mask "pfctx" and "pfcrx" in the `MLX4_EN` module must be set to a value of "0". For example :

```
HostM0 # cat /sys/module/mlx4_en/parameters/pfctx
0
```

```
HostM0 # cat /sys/module/mlx4_en/parameters/pfcrx
0
```

If either or both priority bit masks is set to any other value, they can be set to 0 using either of the following commands:

For SuSE:

```
echo "options mlx4_en pfctx=0 pfcrx=0" >> /etc/modprobe.conf.local
service openibd restart
```

For RHEL:

```
echo "options mlx4_en pfctx=0 pfcrx=0" >> /etc/modprobe.d/modprobe.conf
service rdma restart
```

7. Edit the network configuration files to configure a static IP address for each communication adapter port. The following file listings show the network adapter configuration for the CFs, hosts *cf1* and *cf2*, and members, *member1*, *member2*, *member3*, and *member4*. Edit the network configuration files on each host so that the first communication adapter port listed on each host is on the same subnet as the other hosts. If configuring multiple communication adapter ports on the CFs, pair the additional communication adapter ports CFs so that each DEVICE on the secondary CF is on the same subnetwork as the DEVICE on the primary with the same ID.

```
ssh cf1 cat /etc/sysconfig/network/ifcfg-eth0
DEVICE=eth0
HWADDR=00:02:C9:10:F7:26
TYPE=Ethernet
IPADDR='192.168.1.227'
NETMASK='255.255.255.0'
MTU=''
NAME='Mellanox MT26448 [ConnectX EN 10GigE, PCIe 2.0 5GT/s]'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

```
ssh cf1 cat /etc/sysconfig/network-scripts/ifcfg-eth1
DEVICE=eth1
HWADDR=00:02:C9:10:F7:26
TYPE=Ethernet
IPADDR='192.168.3.227'
NETMASK='255.255.255.0'
MTU=''
NAME='Mellanox MT26448 [ConnectX EN 10GigE, PCIe 2.0 5GT/s]'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

```
ssh cf1 cat /etc/sysconfig/network-scripts/ifcfg-eth2
DEVICE=eth2
HWADDR=00:02:C9:10:F7:26
TYPE=Ethernet
IPADDR='192.168.2.227'
NETMASK='255.255.255.0'
MTU=''
NAME='Mellanox MT26448 [ConnectX EN 10GigE, PCIe 2.0 5GT/s]'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

```
ssh cf1 cat /etc/sysconfig/network-scripts/ifcfg-eth3
DEVICE=eth3
HWADDR=00:02:C9:10:F7:26
TYPE=Ethernet
IPADDR='192.168.4.227'
NETMASK='255.255.255.0'
MTU=''
NAME='Mellanox MT26448 [ConnectX EN 10GigE, PCIe 2.0 5GT/s]'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

```
ssh cf2 cat /etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE=eth0
HWADDR=00:02:C9:10:F7:26
TYPE=Ethernet
IPADDR='192.168.1.228'
```

```

NETMASK='255.255.255.0'
MTU=''
NAME='Mellanox MT26448 [ConnectX EN 10GigE, PCIe 2.0 5GT/s]'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'

```

```

ssh cf2 cat /etc/sysconfig/network-scripts/ifcfg-eth1
DEVICE=eth1
HWADDR=00:02:C9:10:F7:26
TYPE=Ethernet
IPADDR='192.168.3.228'
NETMASK='255.255.255.0'
MTU=''
NAME='Mellanox MT26448 [ConnectX EN 10GigE, PCIe 2.0 5GT/s]'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'

```

```

ssh cf2 cat /etc/sysconfig/network-scripts/ifcfg-eth2
DEVICE=eth2
HWADDR=00:02:C9:10:F7:26
TYPE=Ethernet
IPADDR='192.168.2.228'
NETMASK='255.255.255.0'
MTU=''
NAME='Mellanox MT26448 [ConnectX EN 10GigE, PCIe 2.0 5GT/s]'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'

```

```

ssh cf2 cat /etc/sysconfig/network-scripts/ifcfg-eth3
DEVICE=eth3
HWADDR=00:02:C9:10:F7:26
TYPE=Ethernet
IPADDR='192.168.4.228'
NETMASK='255.255.255.0'
MTU=''
NAME='Mellanox MT26448 [ConnectX EN 10GigE, PCIe 2.0 5GT/s]'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'

```

```

ssh member1 cat /etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE=eth0
HWADDR=00:02:C9:10:F7:26
TYPE=Ethernet
IPADDR='192.168.1.225'
NETMASK='255.255.255.0'
MTU=''
NAME='Mellanox MT26448 [ConnectX EN 10GigE, PCIe 2.0 5GT/s]'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'

```

```

ssh member2 cat /etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE=eth0
HWADDR=00:02:C9:10:F7:26
TYPE=Ethernet
IPADDR='192.168.1.226'
NETMASK='255.255.255.0'
MTU=''

```



```
NAME='Mellanox MT26448 [ConnectX EN 10GigE, PCIe 2.0 5GT/s]'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

```
ssh member3 cat /etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE=eth0
HWADDR=00:02:C9:10:F7:26
TYPE=Ethernet
IPADDR='192.168.1.229'
NETMASK='255.255.255.0'
MTU=''
NAME='Mellanox MT26448 [ConnectX EN 10GigE, PCIe 2.0 5GT/s]'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

```
ssh member4 cat /etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE=eth0
HWADDR=00:02:C9:10:F7:26
TYPE=Ethernet
IPADDR='192.168.1.230'
NETMASK='255.255.255.0'
MTU=''
NAME='Mellanox MT26448 [ConnectX EN 10GigE, PCIe 2.0 5GT/s]'
NETWORK=''
REMOTE_IPADDR=''
STARTMODE='auto'
USERCONTROL='no'
```

Note:

- For simplicity, the IP addresses used in the previous example use the 255.255.255.0 subnetwork mask (NETMASK) so that the subnetwork third and forth IP segments can match the numbers of the interface devices and hostname. This subnetwork mask results in the IP addresses for CFs formatted like 10.222.*interface-id-device-number.CF-hostname-suffix* and members IP addresses like 10.222.*interface-id-device-number.10member-hostname-suffix*.
 - The first communication adapter port on each CF host is on the same subnet as the members.
 - Each communication adapter port on a CF is on a distinct subnet.
 - Communication adapter ports with the same interface *DEVICE* name on the primary and secondary CFs share the same subnet.
8. If configuring multiple communication adapter ports on members, use the same IP subnet for each adapter interface device on the second host as was used for adapter interface with the same device ID on the other hosts so that matching devices are on the same IP subnets.

```
cat /etc/sysconfig/network/ifcfg-ib0
DEVICE=ib0
BOOTPROTO='static'
IPADDR='10.1.1.161'
NETMASK='255.255.255.0'
STARTMODE='onboot'
WIRELESS='no'
```

```
cat /etc/sysconfig/network/ifcfg-ib1
DEVICE=ib1
BOOTPROTO='static'
```

```
IPADDR='10.1.2.161'  
NETMASK='255.255.255.0'  
STARTMODE='onboot'  
WIRELESS='no'
```

All members must be on an IP subnet used by the CF adapter interfaces. The resulting IP subnets are:

- The 10.1.1 subnet has the ib0 device from all members and all CFs.
 - The 10.1.2 subnet has the ib1 device from all members and all CFs.
9. Set up the IP interfaces on the switch. For more information, see “Setting up the IP interfaces on the switch on a 10GE network (Linux)” on page 216.
 10. For all switches in the cluster, disable the Converged Enhance Ethernet (CEE) feature and ensure Global Pause (IEEE 802.3x) is enabled. For a BNT switch with firmware level 6.8.2 and higher, port flow control must also be enabled for Global Pause. For instruction, refer to the switch manual.
 11. Set up netmon.cf on each host. For more information, see “Setting up the netmon.cf file on a 10GE network (Linux).”
 12. Update the /etc/hosts file on each of the hosts so that for each host in the planned DB2 pureScale environment, the file includes all the IP addresses of all the communication adapter ports for all hosts in the planned environment.

For example, in a planned DB2 pureScale environment with multiple communication adapter ports on the CFs and four members, the /etc/hosts configuration file might resemble the following file:

```
192.168.1.227 cf1-eth1.torolab.ibm.com cf1-eth1  
192.168.3.227 cf1-eth2.torolab.ibm.com cf1-eth2  
192.168.2.227 cf1-eth3.torolab.ibm.com cf1-eth3  
192.168.4.227 cf1-eth4.torolab.ibm.com cf1-eth4  
192.168.1.228 cf2-eth1.torolab.ibm.com cf2-eth1  
192.168.3.228 cf2-eth2.torolab.ibm.com cf2-eth2  
192.168.2.228 cf2-eth3.torolab.ibm.com cf2-eth3  
192.168.4.228 cf2-eth4.torolab.ibm.com cf2-eth4  
192.168.1.225 member0-eth1.torolab.ibm.com member0-eth1  
192.168.2.225 member0-eth2.torolab.ibm.com member0-eth2  
192.168.1.226 member1-eth1.torolab.ibm.com member1-eth1  
192.168.2.226 member1-eth2.torolab.ibm.com member1-eth2  
192.168.1.229 member2-eth1.torolab.ibm.com member2-eth1  
192.168.2.229 member2-eth2.torolab.ibm.com member2-eth2  
192.168.1.230 member3-eth1.torolab.ibm.com member3-eth1  
192.168.2.230 member3-eth2.torolab.ibm.com member3-eth2
```

Note:

- In a four member environment that uses only one communication adapter port for each CF and member, the file would look similar to the previous example, but contain only the first IP address of each of the CFs or members.
13. Restart the service for the RoCE subsystem.
service rdma restart

What to do next

Modify the kernel parameters of hosts that you plan to include in the DB2 pureScale environment.

Setting up the netmon.cf file on a 10GE network (Linux):

On a 10 Gigabit Ethernet (10GE) network, one or more pingable IP addresses must be manually set up in the netmon.cf configuration file.

Before you begin

The examples in this topic are based on the figure at the end of this topic, Two CFs and four members connect to two switches.

Procedure

To set up the `netmon.cf` configuration file:

1. Login to the host as root.
2. Retrieve the cluster manager domain name.
`/home/instname/sqllib/bin/db2cluster -cm -list -domain`
3. Stop the domain.
`/home/instname/sqllib/bin/db2cluster -cm -stop -domain domainname -force`
4. Determine which IP address should be entered into the members' `netmon.cf` configuration file. On the member host, to check the communication adapter ports and the associated destination IP subnet, run the **route** command.
`/sbin/route | grep -v link-local`

For example, based on the figure at the end of this topic:

```
Member 0
[root@host3]# route | grep -v link-local
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
192.168.1.0 * 255.255.255.0 U 0 0 0 eth0
192.168.2.0 * 255.255.255.0 U 0 0 0 eth1
9.26.92.0 * 255.255.254.0 U 0 0 0 eth2
default 9.26.92.1 0.0.0.0 UG 0 0 0 eth2
```

```
Member 2
[root@host5]# route | grep -v link-local
Kernel IP routing table
Destination Gateway Genmask Flags Metric Ref Use Iface
192.168.1.0 * 255.255.255.0 U 0 0 0 eth0
192.168.2.0 * 255.255.255.0 U 0 0 0 eth1
9.26.92.0 * 255.255.254.0 U 0 0 0 eth2
default 9.26.92.1 0.0.0.0 UG 0 0 0 eth2
```

The last column (with column name "Iface") lists the adapters on the current host. Choose the adapter that corresponds to the target communication adapter port. In this example, "eth0" and "eth1" are the target 10GE adapters. The corresponding IP addresses in the first column shows the target IP subnet to be used in the next step. In this case, the IP subnets are "192.168.1.0" and "192.168.2.0".

5. With the IP subnet, use the IP interfaces created on the switch 1 and switch 2 that the current host connects to with the same IP subnet. (The IP interface should already be created as part of the 10GE network configuration steps, for details see "Setting up the IP interfaces on the switch on a 10GE network (Linux)" on page 216.) In this example, assuming the IP interfaces on switch 1 have IP addresses of 192.168.1.2 and 192.168.2.2, and switch 2 have IP addresses of 192.168.1.5 and 192.168.2.5, these entries are added to the members configuration file `/var/ct/cfg/netmon.cf`.

```
Member0 (host3)
!REQD eth0 192.168.1.2
!REQD eth1 192.168.2.5
```

```
Member2 (host5)
!REQD eth0 192.168.1.5
!REQD eth1 192.168.2.2
```

where:

- token1 - !REQD is required entity
- token2 - eth0 and eth1 are the RoCE adapter interface names on the local host
- token3 - 192.168.1.2, 192.168.2.5, 192.168.1.5, and 192.168.2.2 are the external pingable IP addresses assigned to the interface created on the switches

The following is an example of what the full configuration file `/var/ct/cfg/netmon.cf` looks like for members:

```
Member0(host3)
!IBQPORTONLY !ALL
!REQD eth2 9.26.92.1
!REQD eth0 192.168.1.2
!REQD eth1 192.168.2.5
!REQD eth0 192.168.1.5
!REQD eth1 192.168.2.2
```

```
Member2(host5)
!IBQPORTONLY !ALL
!REQD eth2 9.26.92.1
!REQD eth0 192.168.1.2
!REQD eth1 192.168.2.5
!REQD eth0 192.168.1.5
!REQD eth1 192.168.2.2
```

6. Determine which IP address should be entered into the cluster caching facilities (CFs) `netmon.cf` configuration file. To check the communication adapter port and the associated destination IP subnet, enter:

```
/sbin/route | grep -v link-local
```

For example:

```
Host1> $ /sbin/route | grep -v link-local
```

Kernel IP routing table

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
192.168.4.0	*	255.255.255.0	U	0	0	0	eth3
192.168.3.0	*	255.255.255.0	U	0	0	0	eth1
192.168.2.0	*	255.255.255.0	U	0	0	0	eth2
192.168.1.0	*	255.255.255.0	U	0	0	0	eth0
9.26.92.0	*	255.255.252.0	U	0	0	0	eth2
default	rsb-v94-hsrp.to	0.0.0.0	UG	0	0	0	eth2

The last column (Iface) indicates the adapter interface name. In this case, eth0, eth1, eth2, and eth3 are the only communication adapter port interface on this host. Four IP subnets are relevant to this host.

All four IP addresses created on the switch (which covers all four IP subnets) must be entered into this host's `netmon.cf` configuration file. For example:

```
!IBQPORTONLY !ALL
!REQD eth2 9.26.92.1
!REQD eth0 192.168.1.2
!REQD eth1 192.168.3.2
!REQD eth7 192.168.2.2
!REQD eth6 192.168.4.2
```

Repeat this step for the secondary CF host in the cluster.

7. Restart the domain.

```
home/instname/sql/lib/bin/db2cluster -cm -start -domain domainname
```

8. Verify all adapters are stable by running the `lssrc` command:

```
lssrc -ls cthats
```

The output is similar to the following:

```

[root@coralm234 ~]# lssrc -ls cthats
Subsystem      Group      PID      Status
cthats         cthats     31938    active
Network Name   Indx Defd  Mbrs  St  Adapter ID      Group ID
CG1             [ 0] 3    3    S   192.168.1.234   192.168.1.234
CG1             [ 0] eth0      0x46d837fd      0x46d83801
HB Interval = 0.800 secs. Sensitivity = 4 missed beats
Ping Grace Period Interval = 60.000 secs.
Missed HBs: Total: 0 Current group: 0
Packets sent    : 560419 ICMP 0 Errors: 0 No mbuf: 0
Packets received: 537974 ICMP 0 Dropped: 0
NIM's PID: 31985
CG2             [ 1] 4    4    S   9.26.93.226     9.26.93.227
CG2             [ 1] eth2      0x56d837fc      0x56d83802
HB Interval = 0.800 secs. Sensitivity = 4 missed beats
Ping Grace Period Interval = 60.000 secs.
Missed HBs: Total: 0 Current group: 0
Packets sent    : 515550 ICMP 0 Errors: 0 No mbuf: 0
Packets received: 615159 ICMP 0 Dropped: 0
NIM's PID: 31988
CG3             [ 2] 3    3    S   192.168.3.234   192.168.3.234
CG3             [ 2] eth1      0x46d837fe      0x46d83802
HB Interval = 0.800 secs. Sensitivity = 4 missed beats
Ping Grace Period Interval = 60.000 secs.
Missed HBs: Total: 0 Current group: 0
Packets sent    : 493188 ICMP 0 Errors: 0 No mbuf: 0
Packets received: 537949 ICMP 0 Dropped: 0
NIM's PID: 31991
CG4             [ 3] 2    2    S   192.168.2.234   192.168.2.234
CG4             [ 3] eth6      0x46d83800      0x46d83803
HB Interval = 0.800 secs. Sensitivity = 4 missed beats
Ping Grace Period Interval = 60.000 secs.
Missed HBs: Total: 0 Current group: 0
Packets sent    : 470746 ICMP 0 Errors: 0 No mbuf: 0
Packets received: 537992 ICMP 0 Dropped: 0
NIM's PID: 31994
CG5             [ 4] 2    2    S   192.168.4.234   192.168.4.234
CG5             [ 4] eth7      0x46d837ff      0x46d83804
HB Interval = 0.800 secs. Sensitivity = 4 missed beats
Ping Grace Period Interval = 60.000 secs.
Missed HBs: Total: 0 Current group: 0
Packets sent    : 470750 ICMP 0 Errors: 0 No mbuf: 0
Packets received: 538001 ICMP 0 Dropped: 0
NIM's PID: 31997
  2 locally connected Clients with PIDs:
rmcd( 32162) hagsd( 32035)
  Dead Man Switch Enabled:
    reset interval = 1 seconds
    trip interval = 67 seconds
  Watchdog module in use: softdog
  Client Heartbeating Enabled. Period: 6 secs. Timeout: 13 secs.
  Configuration Instance = 1322793087
  Daemon employs no security
  Segments pinned: Text Data Stack.
  Text segment size: 650 KB. Static data segment size: 1475 KB.
  Dynamic data segment size: 2810. Number of outstanding malloc: 1165
  User time 32 sec. System time 26 sec.
  Number of page faults: 0. Process swapped out 0 times.
  Number of nodes up: 4. Number of nodes down: 0.

```

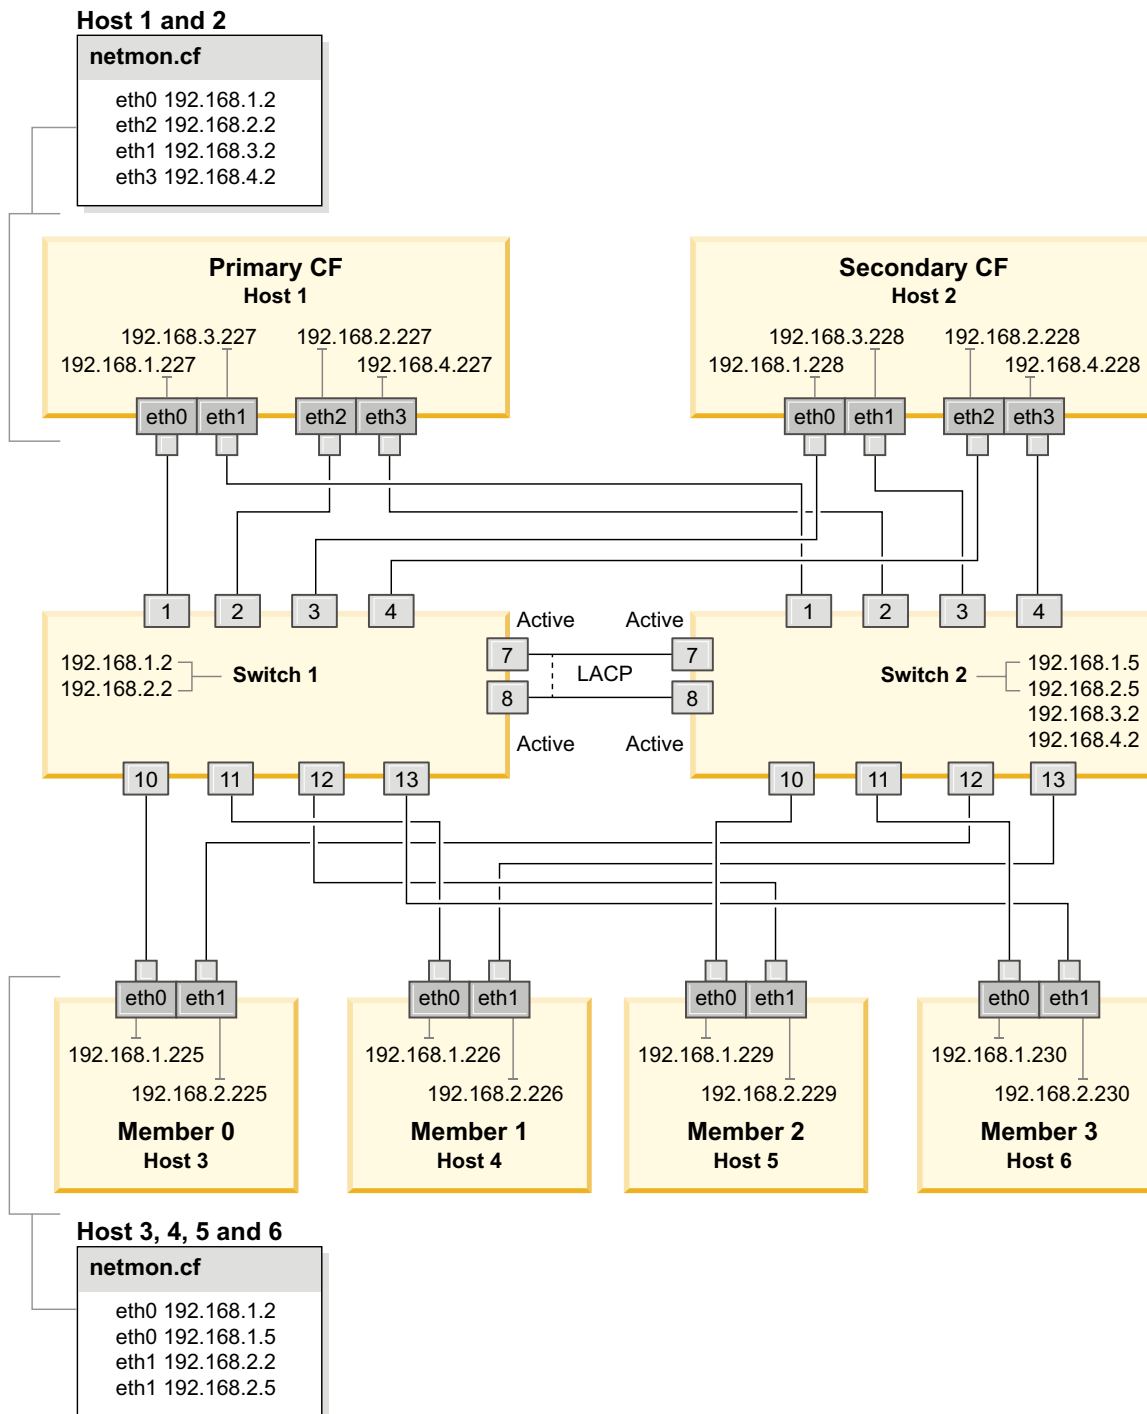


Figure 23. Two CFs and four members connect to two switches.

Modifying kernel parameters (Linux)

The database manager uses a formula to automatically adjust kernel parameter settings and eliminate the need for manual updates to these settings.

Before you begin

You must have root authority to modify kernel parameters.

Procedure

To update kernel parameters on Red Hat and SUSE Linux:

1. Run the **ipcs -l** command to list the current kernel parameter settings.
2. Analyze the command output to determine whether you have to change kernel settings or not by comparing the current values with the *enforced minimum settings* at <http://publib.boulder.ibm.com/infocenter/db2luw/v9r8/topic/com.ibm.db2.luw.qb.server.doc/doc/c0057140.html>. The following text is an example of the **ipcs** command output with comments added after **//** to show what the parameter names are:

```
# ipcs -l

----- Shared Memory Limits -----
max number of segments = 4096           // SHMMNI
max seg size (kbytes) = 32768           // SHMMAX
max total shared memory (kbytes) = 8388608 // SHMALL
min seg size (bytes) = 1

----- Semaphore Limits -----
max number of arrays = 1024             // SEMMNI
max semaphores per array = 250          // SEMMSL
max semaphores system wide = 256000    // SEMMNS
max ops per semop call = 32             // SEMOPM
semaphore max value = 32767

----- Messages: Limits -----
max queues system wide = 1024           // MSGMNI
max size of message (bytes) = 65536     // MSGMAX
default max size of queue (bytes) = 65536 // MSGMNB
```

- Beginning with the first section on Shared Memory Limits, the **SHMMAX** limit is the maximum size of a shared memory segment on a Linux system. The **SHMALL** limit is the maximum allocation of shared memory pages on a system.
 - It is recommended to set the **SHMMAX** value to be equal to the amount of physical memory on your system. However, the minimum required on x86 systems is 268435456 (256 MB) and for 64-bit systems, it is 1073741824 (1 GB).
 - The **SHMALL** parameter is set to 8 GB by default (8388608 KB = 8 GB). If you have more physical memory than 8 GB, and it is to be used for DB2, then this parameter increases to approximately 90% of your computer's physical memory. For instance, if you have a computer system with 16 GB of memory to be used primarily for DB2, then **SHMALL** should be set to 3774873 (90% of 16 GB is 14.4 GB; 14.4 GB is then divided by 4 KB, which is the base page size). The **ipcs** output converted **SHMALL** into kilobytes. The kernel requires this value as a number of pages. If you are upgrading to DB2 Version 10.1 and you are not using the default **SHMALL** setting, you must increase the **SHMALL** setting by an additional 4 GB. This increase in memory is required by the fast communication manager (FCM) for additional buffers or channels.
- The next section covers the amount of semaphores available to the operating system. The kernel parameter **sem** consists of four tokens, **SEMMSL**, **SEMMNS**, **SEMOPM** and **SEMMNI**. **SEMMNS** is the result of **SEMMSL** multiplied by **SEMMNI**. The database manager requires that the number of arrays (**SEMMNI**) be increased as necessary. Typically, **SEMMNI** should be twice the maximum number of agents expected on the system multiplied by the number of logical partitions on the database server computer plus the number of local application connections on the database server computer.

- The third section covers messages on the system.
 - The **MSGMNI** parameter affects the number of agents that can be started; the **MSGMAX** parameter affects the size of the message that can be sent in a queue, and the **MSGMNB** parameter affects the size of the queue.
 - The **MSGMAX** parameter should be changed to 64 KB (that is, 65536 bytes), and the **MSGMNB** parameter should be increased to 65536.
- 3. Modify the kernel parameters that you have to adjust by editing the `/etc/sysctl.conf` file. If this file does not exist, create it. The following lines are examples of what should be placed into the file:


```
#Example for a computer with 16GB of RAM:
kernel.shmmni=4096
kernel.shmmax=17179869184
kernel.shmall=8388608
#kernel.sem=<SEMMSL> <SEMMNS> <SEMOPM> <SEMMNI>
kernel.sem=250 1024000 32 4096
kernel.msgmni=16384
kernel.msgmax=65536
kernel.msgmnb=65536
```
- 4. Run **sysctl** with **-p** parameter to load in **sysctl** settings from the default file `/etc/sysctl.conf`:


```
sysctl -p
```
- 5. Optional: Have the changes persist after every reboot:
 - (SUSE Linux) Make `boot.sysctl` active.
 - (Red Hat) The `rc.sysinit` initialization script reads the `/etc/sysctl.conf` file automatically.

For the latest information about supported Linux distributions, see <http://www.ibm.com/software/data/db2/linux/validate/>.

Preinstallation checklist for DB2 pureScale Feature (Linux)

Perform the following preinstallation steps and verify them on each host before installing the IBM DB2 pureScale Feature.

Before you install

The following steps must be performed on all hosts:

1. All hosts must use the same Linux distribution.
2. DB2 pureScale instances require specific users and groups, including fenced users. You can create the users before starting the DB2 Setup wizard or have the wizard create them for you as you progress through the panels. If you are not creating or modifying instances you can create the required users after completing the installation.
3. Ensure that the required Linux Version and Service Pack is installed.
 - SUSE Linux Enterprise Server (SLES) 10 Service Pack (SP) 3 - The minimum required level is the x64 version of SUSE SLES 10 SP3, kernel 2.6.16.60-0.69.1-smp and the matching kernel source. Check the `/etc/SuSE-release` file for the operating system level and service pack. The following sample output should be returned:

```
cat /etc/SuSE-release
SUSE Linux Enterprise Server 10 (x86_64)
VERSION = 10
PATCHLEVEL = 3
```

Enter the following command:


```
cat /proc/version
Linux version 2.6.16.60-0.69.1-smp (geeko@buildhost)
(gcc version 4.1.2 20070115 (SUSE Linux)) #1 SMP Fri May 28 12:10:21 UTC 2010
```

- For single InfiniBand communication adapter port on Red Hat Enterprise Linux (RHEL) 5.6 - The minimum required level is the x64 version of RHEL 5.6 and the matching kernel source. Check the `/etc/redhat-release` file for the operating system level and service pack. The following sample output should be returned for RHEL 6.1:

```
cat /etc/redhat-release
Red Hat Enterprise Linux Server release 6.1 (Santiago)
```

and for RHEL 5.7:

```
cat /etc/redhat-release
Red Hat Enterprise Linux Server release 5.7 (Tikanga)
```

Enter the following command for RHEL 6.1:

```
cat /proc/versionLinux version
2.6.32-131.0.15.el6.x86_64 (mockbuild@x86-007.build.bos.redhat.com)
(gcc version 4.4.4 20100726 (Red Hat 4.4.4-13) (GCC) )
#1 SMP Tue May 10 15:42:40 EDT 2011
```

and for RHEL 5.7:

```
cat /proc/versionLinux version
2.6.18-274.7.1.el5 (mockbuild@x86-004.build.bos.redhat.com) (gcc version
4.1.2 20080704 (Red Hat 4.1.2-50)) #1 SMP Mon Oct 17 11:57:14 EDT 2011
```

Note: If kernel modules (such as RDAC) have not been rebuilt after the kernel upgrade, the kernel modules must be rebuilt before proceeding.

- For multiple InfiniBand communication adapter port and single or multiple 10GE communication adapter port, Red Hat Enterprise Linux (RHEL) 6.1 is required. - The minimum required level is the x64 version of RHEL 6.1, kernel and the matching kernel source. Check the `/etc/redhat-release` file for the operating system level and service pack. The following sample output is returned:

```
cat /etc/redhat-release
Red Hat Enterprise Linux Server release 6.1 (Santiago)
```

Enter the following command:

```
cat /proc/version
Linux version
```

Ensure the following 32-bit RSCT packages are installed:

- libibcm.i686
- libibverbs-rocee.i686
- librdmacm.i686
- libcxgb3.i686
- libibmad.i686
- libibumad.i686
- libmlx4-rocee.i686
- libmthca.i686

As root, run the following command on each of the package names (listed above):

```
yum list | grep package_name
```

For example:

```
[root]# for i in `cat /tmp/list`; do yum list | grep $i; done
libibcm.i686          1.0.5-2.el6      @rhel-x86_64-server-6
libibverbs-rocee.i686 1.1.4-4.el6      @rhel-x86_64-server-hpn-6
librdmacm.i686        1.0.10-2.el6     @rhel-x86_64-server-6
libcxb3.i686          1.3.0-1.el6      @rhel-x86_64-server-6
libibmad.i686         1.3.4-1.el6      @rhel-x86_64-server-6
libibumad.i686        1.3.4-1.el6      @rhel-x86_64-server-6
libmlx4-rocee.i686    1.0.1-8.el6      @rhel-x86_64-server-hpn-6
libmthca.i686         1.0.5-7.el6      @rhel-x86_64-server-6
```

4. For single and multiple communication adapter ports on InfiniBand network on SLES and single communication adapter port on InfiniBand network on RHEL 5.6, ensure that OpenFabrics Enterprise Distribution (OFED) software is installed, and configured. See Configuring the networking settings of hosts on a 10GE network (Linux) and Configuring the networking settings of hosts on an InfiniBand network (Linux) for more information.
5. Ensure that OpenSSH is installed from the SLES10 media or RHEL media, as appropriate.
6. For InfiniBand network on both SLES and RHEL 5.5 and 10GE network on RHEL 5.5, ensure that the **openibd** service is enabled.

```
# chkconfig --list | grep -i openibd
openibd 0:off 1:off 2:off 3:off 4:off 5:off 6:off
```

The **openibd** service must be enabled. To enable the service:

```
# chkconfig openibd on
# chkconfig --list | grep -i openibd
openibd 0:off 1:off 2:on 3:on 4:on 5:on 6:off
```

For 10GE network on RHEL 6.1, ensure that the **rdma** service is enabled:

```
chkconfig --list | grep -i rdma
rdma 0:off 1:off 2:off 3:off 4:off 5:off 6:off
```

The **rdma** service must be enabled. To enable the service:

```
# chkconfig rdma on
# chkconfig --list | grep -i rdma
rdma 0:on 1:on 2:on 3:on 4:off 5:off 6:off
```

7. DB2 pureScale Feature requires libstdc++.so.6. Verify that the files exist with the following commands:

```
ls /usr/lib/libstdc++.so.6*
ls /usr/lib64/libstdc++.so.6*
```

8. Optional. To use a specific set of ports, ensure that the ports are free on all hosts. Otherwise, the installer selects the unique set of ports across all hosts. The Fast Communications Manager (FCM) requires a port range of the three mandatory ports plus the value provided for the logical members field. This port range can designate up to 130 hosts (128 members + 2 cluster caching facilities.) The default FCM start port is 60000 and must be in the range 1024 - 65535. In addition, two ports are required for cluster caching facilities. These two ports are chosen automatically.

Use the **grep** command on the `/etc/services` file to ensure that a contiguous range of ports is available.

9. Confirm that the required network adapters are installed on each server. Ensure that an Ethernet network (*eth0*) and an InfiniBand network (*ib0*) or 10 Gigabit Ethernet network (*eth1*) display. The following sample uses the **netstat -i** command to list all available network adapters with an InfiniBand communication adapter port.

```

root@host1: /> netstat -i
Iface  MTU Met    RX-OK RX-ERR RX-DRP RX-OVR   TX-OK TX-ERR TX-DRP TX-OVR Flg
eth0    1500  0 6876034      0      0      0 5763121      0      0      0 BMRU
ib0     65520  0 106972       0      0      0      9      0      0      0 BMRU
lo      16436  0 180554       0      0      0 180554      0      0      0 LRU

```

Note: The DB2 pureScale Feature does not support a mixed environment of InfiniBand and 10 Gigabit Ethernet networks, all servers must use the same communication adapter port.

10. As root, validate ssh access between all hosts. From the current host, run the **hostname** command on the current host and on all other hosts in the cluster by using the **ssh** command. The result of the **hostname** command matching the host name identified in the **ssh** command verifies ssh access between the two hosts.

```

$ ssh host1 hostname
host1

```

11. Optional. For DB2 managed GPFS installations, verify the remote shell and remote file copy settings default to **db2locssh** and **db2scp**. For example:

```

/usr/lpp/mmfs/bin/mmlscluster
Remote shell command:    /var/db2/db2ssh/db2locssh
Remote file copy command: /var/db2/db2ssh/db2scp

```

12. If upgrading from DB2 Version 9.8 Fix Pack 2 or earlier, ensure that the .update file, located at <db2 instance shared directory>/sqlllib_shared/.update, is synchronized correctly after adding or dropping a member or cluster caching facility (CF). An example of the file location is /db2sd_20110126085343/db2sdin1/sqlllib_shared/.update, where <db2 instance shared directory>=db2sd_20110126085343.

To ensure correct synchronization, check that all hosts are listed in the .update file with the following format: *hostname=install path*. If incorrectly formatted, update the file. For example: *machineA=/opt/IBM/db2/V9.8*, where *hostname=machineA* and *install path=/opt/IBM/db2/V9.8*.

13. As root, ensure that the /tmp directory has at least 5 GB of free space. The following command shows the free space in the /tmp directory.

```

$ cd /tmp
$ df -k .
Filesystem 1K-blocks Used Available Use% Mounted on
/dev/sda7 2035606 164768 1870838 9% /tmp

```

14. Determine the number of paths to your device on the system with IBM RDAC, DM-MP, or EMC PowerPath driver:

On systems with IBM RDAC driver, the following are the commands to run, and sample output:

- a. Determine the LUN mapping by using the **lsvdev** command:

```

host1:~ # /opt/mpp/lsvdev
Array Name      Lun      sd device
-----
DS5300SVT1      0        -> /dev/sdc
DS5300SVT1      1        -> /dev/sdd
DS5300SVT1      2        -> /dev/sde
DS5300SVT1      3        -> /dev/sdf
DS5300SVT1      4        -> /dev/sdg

```

- b. Get a list of storage arrays seen by the host:

```

host1:~ # /usr/sbin/mppUtil -a
Hostname       = host1
Domainname     = N/A
Time           = GMT 08/06/2010 16:27:59

```

ID	WWN	Type	Name
0	600a0b800012abc6000000000402756fc	FC	FASTSVT1
1	600a0b800047bf3c0000000004a9553b8	FC	DS5300SVT1

- ```
host1:~ # /usr/sbin/mppUtil -a DS5300SVTI | awk '/Status/ || /NumberOfPaths/'
Controller 'A' Status:
 NumberOfPaths: 1 FailoverInProg: N
Controller 'B' Status:
 NumberOfPaths: 1 FailoverInProg: N
```

a. Check the Linux SCSI devices:

```
[root@host1 ~]# lsscsi
[3:0:1:0] disk IBM 2107900 .450 /dev/sdk
[3:0:1:2] disk IBM 2107900 .450 /dev/sdl
[3:0:1:3] disk IBM 2107900 .450 /dev/sdm
[3:0:1:4] disk IBM 2107900 .450 /dev/sdn
[3:0:2:0] disk IBM 2107900 .450 /dev/sdo
[3:0:2:2] disk IBM 2107900 .450 /dev/sdp
[3:0:2:3] disk IBM 2107900 .450 /dev/sdq
[3:0:2:4] disk IBM 2107900 .450 /dev/sdr
[4:0:0:0] disk IBM 2107900 .450 /dev/sdc
[4:0:0:2] disk IBM 2107900 .450 /dev/sdd
[4:0:0:3] disk IBM 2107900 .450 /dev/sde
[4:0:0:4] disk IBM 2107900 .450 /dev/sdf
[4:0:1:0] disk IBM 2107900 .450 /dev/sdg
[4:0:1:2] disk IBM 2107900 .450 /dev/sdh
[4:0:1:3] disk IBM 2107900 .450 /dev/sdi
[4:0:1:4] disk IBM 2107900 .450 /dev/sdj
```

- ```
[root@host1 ~]# sg_map -x
/dev/sg9    4 0 0 0 0 /dev/sdc
/dev/sg10   4 0 0 2 0 /dev/sdd
/dev/sg11   4 0 0 3 0 /dev/sde
/dev/sg12   4 0 0 4 0 /dev/sdf
/dev/sg13   4 0 1 0 0 /dev/sdg
/dev/sg14   4 0 1 2 0 /dev/sdh
/dev/sg15   4 0 1 3 0 /dev/sdi
/dev/sg16   4 0 1 4 0 /dev/sdj
/dev/sg17   3 0 1 0 0 /dev/sdk
/dev/sg18   3 0 1 2 0 /dev/sdl
/dev/sg19   3 0 1 3 0 /dev/sdm
/dev/sg20   3 0 1 4 0 /dev/sdn
/dev/sg21   3 0 2 0 0 /dev/sdo
/dev/sg22   3 0 2 2 0 /dev/sdp
/dev/sg23   3 0 2 3 0 /dev/sdq
/dev/sg24   3 0 2 4 0 /dev/sdr
```

- d.

238 Installing DB2 Servers

```

|
|----- user friendly
| name
[size=100G] [features=1 queue_if_no_path] [hwhandler=0] [rw]
|
|----- hardware handler,
| if any (seen in
| cases of FastT, EMC)
|
|----- features supported
| or configured
|----- Size of the disk

\_ round-robin 0 [prio=0] [active]
|
|----- Path Group State
|----- Path Group Priority
|
|----- Path Selector and
| repeat count
|----- Path Group Level

\_ 4:0:0:0 sdc 8:32 [active] [ready]
|
|----- Physical Path State
|----- Device Mapper State
|----- Major, Minor number
| of disk
|----- Linux SCSI device name
|----- SCSI Information:
| Host_ID, Channel_ID,
| SCSI_ID, LUN_ID

\_ 4:0:1:0 sdg 8:96 [active] [ready]
\_ 3:0:1:0 sdk 8:160 [active] [ready]
\_ 3:0:2:0 sdo 8:224 [active] [ready]

mpath6 (36005076304ffc21f0000000000001123) dm-3 IBM,2107900
[size=100G] [features=1 queue_if_no_path] [hwhandler=0] [rw]
\_ round-robin 0 [prio=0] [active]
\_ 4:0:0:4 sdf 8:80 [active] [ready]
\_ 4:0:1:4 sdj 8:144 [active] [ready]
\_ 3:0:1:4 sdn 8:208 [active] [ready]
\_ 3:0:2:4 sdr 65:16 [active] [ready]
mpath5 (36005076304ffc21f0000000000001122) dm-2 IBM,2107900
[size=1.0G] [features=0] [hwhandler=0] [rw]
\_ round-robin 0 [prio=0] [enabled]
\_ 4:0:0:3 sde 8:64 [active] [ready]
\_ 4:0:1:3 sdi 8:128 [active] [ready]
\_ 3:0:1:3 sdm 8:192 [active] [ready]
\_ 3:0:2:3 sdq 65:0 [active] [ready]
mpath4 (36005076304ffc21f0000000000001121) dm-1 IBM,2107900
[size=100G] [features=1 queue_if_no_path] [hwhandler=0] [rw]
\_ round-robin 0 [prio=0] [active]
\_ 4:0:0:2 sdd 8:48 [active] [ready]
\_ 4:0:1:2 sdh 8:112 [active] [ready]
\_ 3:0:1:2 sdl 8:176 [active] [ready]
\_ 3:0:2:2 sdp 8:240 [active] [ready]

```

The block device name is listed as the Linux SCSI device name. If there are multiple paths, multiple block devices are displayed under each pseudo name.

On systems with EMC PowerPath driver, the following are the commands to run, and sample output:

- a. Run the **powermt** command to display all path and device mappings. This command lists the block devices and paths which are mapped to the device path (for example, */dev/emcpowerd* of which the EMC pseudo name is *emcpowerd*):

```

host1:~ # powermt display dev=all
Pseudo name=emcpowerd
Symmetrix ID=000194900547
Logical device ID=0040
state=alive; policy=BasicFailover; priority=0; queued-I/Os=0
=====

```

```

----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-I/Os Errors
=====
      3 qla2xxx          sdg       FA  7eB   active alive      0      0

Pseudo name=emcpowerc
Symmetrix ID=000194900547
Logical device ID=0041
state=alive; policy=BasicFailover; priority=0; queued-I/Os=0
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-I/Os Errors
=====
      3 qla2xxx          sdh       FA  7eB   active alive      0      0

Pseudo name=emcpowerb
Symmetrix ID=000194900547
Logical device ID=0126
state=alive; policy=BasicFailover; priority=0; queued-I/Os=0
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-I/Os Errors
=====
      3 qla2xxx          sdi       FA  7eB   active alive      0      0

Pseudo name=emcpowera
Symmetrix ID=000194900547
Logical device ID=013C
state=alive; policy=BasicFailover; priority=0; queued-I/Os=0
=====
----- Host ----- - Stor - -- I/O Path - -- Stats ---
### HW Path          I/O Paths   Interf.  Mode   State  Q-I/Os Errors
=====
      3 qla2xxx          sdj       FA  7eB   active alive      0      0

```

The block device name is listed under I/O Paths column. If there are multiple paths, multiple block devices are displayed under each pseudo name.

15. Increase the Mellanox HCA driver `mlx4_core` parameter **log_mttts_per_seg** value from 3 (the default) to 7 on the host where the cluster caching facility (CF) resides. To increase the size, issue the following command as root:
 - On SUSE:


```
echo "options mlx4_core log_mttts_per_seg=7" >> /etc/modprobe.conf.local
```
 - On RHEL 6.x:


```
echo "options mlx4_core log_mttts_per_seg=7" >> /etc/modprobe.d/modprobe.conf
options mlx4_core log_mttts_per_seg=7
```

For this change to take effect, you must reboot the server. To check whether your change is effective on the module, issue the following command:

```
<host-name>/sys/module/mlx4_core/parameters # cat /sys/module/mlx4_core/parameters/log_mttts_per_seg
7
```

16. In some installations, if Intel TCO WatchDog Timer Driver modules are loaded by default, they should be blacklisted, so that they do not start automatically or conflict with RSCT. To blacklist the modules, edit the following files:
 - a. To verify if the modules are loaded


```
lsmod | grep -i itco_wdt; lsmod | grep -i itco_vendor_support
```
 - b. Edit the configuration files:
 - On RHEL 5.x and RHEL 6.1, edit file `/etc/modprobe.d/blacklist.conf`:


```
# RSCT hatsd
blacklist itco_wdt
blacklist itco_vendor_support
```
 - On SLES, edit file `/etc/modprobe.d/blacklist`:


```
add
blacklist itco_wdt
blacklist itco_vendor_support
```

Using the DB2 Setup wizard

To install the DB2 pureScale Feature, you must know the following items. You can enter your values for each of these steps in the preinstallation cheat sheet section that follows.

- The name of the directory where you want to install the DB2 pureScale Feature.
- The device path for the shared disk partition where the DB2 database and the shared configuration files of the DB2 pureScale instance are stored by default. You should use the DB2 cluster file system.
- If using a DB2 cluster services tiebreaker, the device path for the tiebreaker partition. This partition must be different from the shared disk partition specified earlier.

The DB2 cluster services tiebreaker is used in a recovery scenario when your DB2 pureScale instance is split in half due to a communication failure. DB2 cluster services use the tiebreaker to determine which half of the DB2 pureScale instance will remain online. This partition must be at least 25 MB in size and accessible from all hosts.

- The host name, short host name, or IP address for each of the hosts you want to participate in the DB2 pureScale environment.
- Optionally, you can specify the names of the hosts you want to set up as cluster caching facilities rather than accept the default choices provided by the DB2 Setup wizard. You can manually select which hosts act as cluster caching facilities, or have the DB2 installer assign them.

Preinstallation cheat sheet

Enter the appropriate required item value in the "Your Value" field.

Table 41. Preinstallation cheat sheet

Required Item	Your Value	Example
Instance owner/group name		<i>db2sdin1/db2iadm1</i>
Fenced user/group name		<i>db2sdfe1/db2fadm1</i>
Installation directory name		<i>/opt/IBM/db2/V10.1</i>
Shared file system disk		<i>/dev/hdisk12</i>
Hosts to include		<i>db2_host01 - db2_host04.</i>
Netname interconnect for each member and CF		InfiniBand network example: <i>db2_<hostname>_ib0</i> 10 Gigabit Ethernet network example: <i>db2_<hostname>_en1</i> Note: <i>db2_<hostname>_en1</i> does not map to a regular ethernet adapter. It must map to the pseudo IP address for the 10GE communication adapter port.

Table 42. Preinstallation cheat sheet - optional items

Optional Item	Your Value	Example
Tiebreaker disk		On AIX: /dev/hdisk13 On Linux: /dev/dm-0 or /dev/sdc
FCM port range		60000 - 60004
cluster caching facilities port range		56000 - 56001
DB2 communication port		50001
Hosts to set up as cluster caching facilities		db2_host03 and db2host04
On InfiniBand, the cluster interconnect netnames of the cluster caching facilities		Primary: db2_<hostname1>_ib0, db2_<hostname1>_ib1, db2_<hostname1>_ib2, db2_<hostname1>_ib3 Secondary: db2_<hostname2>_ib0, db2_<hostname2>_ib1, db2_<hostname2>_ib2, db2_<hostname2>_ib3
On 10GE, the cluster interconnect netnames of the cluster caching facilities		Primary: db2_<hostname1>_en1, db2_<hostname1>_en2, db2_<hostname1>_en3, db2_<hostname1>_en4 Secondary: db2_<hostname2>_en1, db2_<hostname2>_en2, db2_<hostname2>_en3, db2_<hostname2>_en4
Hosts to set up as members		db2_host01 and db2host02

What to do next

If you completed all the steps in the preinstallation checklist and filled out the cheat sheet, you can proceed directly to the installation section.

GDPC infrastructure and GDPC-specific prerequisite conditions

Before setting up a geographically dispersed DB2 pureScale cluster (GDPC), a number of conditions must be met.

Supported operating systems

Geographically dispersed DB2 pureScale cluster (GDPC) is supported on AIX operating systems and RHEL operating systems, with the following networks:

- AIX: InfiniBand (IB) network

- RHEL: 10 Gigabit Ethernet (10GE) network.

For additional operating system prerequisite details, see installation prerequisites for DB2 pureScale Feature.

Site-to-site connectivity


The connection between sites is a key piece of infrastructure in a geographically dispersed DB2 pureScale cluster (GDPC). DB2 pureScale software uses low-latency, high-bandwidth RDMA messaging between members and cluster facilities (CFs), and in a GDPC configuration, many such messages traverse the link from one site to the other.

For example, in the case of an InfiniBand high speed interconnect, Longbow InfiniBand extender technology from Obsidian Strategics provides a transparent connection between the two portions of the high speed interconnect network that are located at the two sites, and maintains the ability to execute RDMA operations across GDPC, even at relatively large distances. Used in pairs at either end of the site-to-site interconnect, the extenders accept a high speed interconnect connection to a site-local high speed interconnect switch, and through it, connect to the members and CF. The extender translates high speed interconnect traffic to and from packets that are sent and received over the site-to-site interconnect (either a 'dark fiber' or 10 GB WAN connection).

The extenders themselves add only a very small amount of extra latency to the DB2 pureScale message protocol. The bulk of the extra site-to-site message latency, when compared to a single-site DB2 pureScale cluster, arises from the simple fact of distance: each kilometer of transmission in glass fiber adds an additional 5 microseconds of delay. So for example, a 10km distance between sites would add $(10\text{km} \times 5 \text{ microseconds/km}) \times 2$ for round trip = 100 microseconds of extra latency for most types of messages. In practical terms, workloads that have higher ratios of read activity compared to write activity tend to trigger fewer message exchanges with the CF, and so would be less impacted by additional site-to-site latency.

Current Longbow IB extenders operate at the 4X width Single Data Rate (SDR) or 10 GB data rate between end points (subject to the capacity of the dark fiber / WAN link.) If redundancy or additional cross-site capacity is required, Longbow units can be 'ganged' in multiple pairs between sites (see Figure 1). As well, different Longbow models provide different features which can be useful in certain circumstances, such as encryption in the E-100 and X-100 models, which might be important if the site-to-site interconnect is shared or public and encryption is required by security policies. All current Longbow models are supported with GDPCs. Particular configurations, such as the choice of model, use of WAN or fiber, or choice of transceiver wavelength, and other characteristics, are not specified here, and should be selected based on the physical infrastructure to be used, IT policies in effect, and so on. For more information about Longbow IB extenders, contact Obsidian Research. (<http://www.obsidianresearch.com/>)

Figure 24. Multiple pairs of Longbow units between sites



Two site or three site configurations

A GDPC is composed of two main sites A and B, with each having an equal number of members and CFs. For example, if site A has two members and one CF, site B must also have two members and one CF. It is a best practice that each main site have the same number of physical machines as well. For example, you do not want to have one machine with four LPARs on one site and two machines with two LPARs each on the other site. One key clustering concept that must be considered is the concept of 'quorum'. Quorum refers to the number of computer systems that must be online in the cluster in order for the cluster to remain operational. There are two types of quorum, operational quorum and configuration quorum. Operational quorum is needed for software services on the cluster to operate. Configuration quorum is needed to apply configuration changes to a cluster, such as adding a new computer system to the cluster. Configuration quorum requires a strict majority of online computer systems in the cluster, so for example in a cluster that comprises 6 computer systems, at least 4 of those computer systems must be online to perform any cluster configuration updates.

In a non-GDPC environment, operational quorum is typically achieved through the use of a tiebreaker disk. In the event of having only half the computer systems in a cluster online (or a network partition where each half is simultaneously online with no network connectivity to the other half), the disk "tiebreaker" device can be acquired by one half of the cluster. This allows it to achieve operational quorum and run software services (that is, the DB2 pureScale instance) on that winning half of the cluster. In the event of a network partition, the "losing" half would consequently be fenced from the cluster, preventing it from accessing any shared instance data. The requirement for a disk tiebreaker, however, is that the single tiebreaker disk must be accessible from each computer system in the cluster. In a GDPC environment, this disk would have to physically reside at one of the two sites, which in the event of a complete network partition between the two sites, would prevent the other site from being able to achieve operational quorum. In the case of clusters with an odd number of nodes, a majority of online nodes is needed for operational quorum. However, in the case where the cluster has an even number of nodes, with an even split of online nodes, a tiebreaker disk decides which subcluster gains operational quorum. In cases where one half of the cluster is down, the online subcluster claims the tiebreaker and gain quorum.

GDPC environments rely on strict majority quorum semantics, where one additional tiebreaker host T is required to maintain quorum in the event of site failure. This tiebreaker host T must be the same architecture type as the machines at the two main sites. For example, it must run the AIX operating system, although

it does not need to be the same hardware model. A best practice is to also be running the same OS level across all computer systems in the cluster. This additional host does not run any DB2 members or CFs.

A two-site configuration where the tiebreaker host is physically located at one of the two main sites would not be able to achieve either operational or configuration quorum in the event of a site failure at the site containing host T. As such, it is a best practice for continuous availability to use a three-site configuration where the tiebreaker host T is physically located at a separate third site (site C), in order to achieve continuous availability in the event of a failure affecting either of the data processing sites (site A or site B), as majority quorum can be established between site C and the surviving data processing site. Note that in three-site configurations, all three sites can be on different IP subnets as long as each computer system from each site is able to “ping” each other computer system in the cluster. Site C also does not require high speed interconnect connectivity; only sites A and B require high speed interconnect connectivity, with a single high speed interconnect subnet spanning both sites. Further, if SCSI-3 PR is not used, site C does not require SAN access, since there are no members or CFs configured to run on host T.

To aid in problem determination, it is a best practice to have all computer systems at all sites configure their system clocks to the same timezone.

Zoned SAN storage

GDPC requires that both sites A and B have direct access to each others’ disks. To this end, a number of options are available for extending a SAN across the data centers. Options include transmitting Fibre Channel (FC) traffic directly over ATM or IP networks, or using iSCSI to transmit SCSI commands over IP. Dark fiber is likely to be the fastest but also the most expensive option.

GPFS synchronous replication

A typical non-GDPC uses GPFS in a non-replicated configuration. In such a case, all GPFS disk activity for a given file system goes to a single GPFS failure group. When disks are not replicated, a disk failure can leave some of the file system data inaccessible. For a GDPC, however, GPFS replication is used between sites A & B in order to ensure that an entire copy of the data is available at the surviving site in the event of a total site failure.

GDPC configuration leverages GPFS replication, by configuring each site to maintain an entire copy of the file system data in its own failure group. As long as quorum is maintained in the cluster, in the event of a site failure (one of the failure groups are lost or inaccessible), the other site can continue with read/write access to the file system.

Tiebreaker host T requires a small disk or partition for each replicated GPFS file system in use by DB2 pureScale software to be used as a file system quorum disk. The amount of storage for each disk or partition is approximately 50 MB, and these disks or partitions only need to be accessible by host T, and are only used to store filesystem descriptors. I/O activity to disks or partitions that are used to store only filesystem descriptors is very low. Using a full physical volume for this purpose is wasteful and not necessarily practical; configuring a small volume is sufficient for this case.

Performance impact over single site

The introduction of significant distances between cluster members at different sites increases DB2 pureScale message latency by an amount of about 5 microseconds per kilometer of glass fiber. In some cases, the amount can be higher, if the connection includes signal repeaters, or is shared with other applications.

Besides distance, the performance overhead experienced by a GDPC configuration also depends on the workloads in use. The greater the portion of write activity (INSERT, UPDATE, DELETE) in the workload, the more messages need to be sent from members to the CFs, and the more disk writes (especially to the transaction logs) need to be made. This typically leads to higher perceived overhead at a given distance. Conversely, a greater portion of read (SELECT) activity means fewer messages and fewer disk writes, and reduced overhead.

DB2 pureScale software is designed to have minimal downtime if a host fails due to hardware or software faults. In the event of a hardware failure, a system must be 'I/O fenced' to prevent it from corrupting the data. A key piece of technology that DB2 pureScale software uses to minimize downtime is SCSI-3 Persistent Reserve (PR). After a host is I/O fenced, it can no longer access the storage device, and any I/O attempt is blocked.

If SCSI-3 PR is not enabled, the GPFS disk lease expiry mechanism is used to fence failed systems. This typically results in a longer recovery time to wait for the lease to expire.

DB2 pureScale Feature installation methods

This topic provides information about the IBM DB2 pureScale Feature installation methods: the DB2 Setup wizard, a response file installation, and manual installation using the **db2_install** command.

DB2 Setup wizard

The DB2 Setup wizard is a GUI installer. The DB2 Setup wizard provides an easy-to-use interface for installing DB2 products and for performing initial setup and configuration tasks.

With the DB2 pureScale Feature release, the DB2 Setup wizard has been changed significantly compared to earlier versions. This change should be of note if you currently use a response file installation as there are new keywords and options.

The DB2 Setup wizard provides a clear flow through which you can launch a DB2 pureScale Feature installation from one host, called the installation-initiating host, and successfully setup a DB2 pureScale environment across multiple hosts. This interactive deployment includes validating all hosts, installing the DB2 pureScale Feature on all hosts, and ensuring that the DB2 clusters services are set up and available between all the hosts. Simply provide a list of host names and two valid shared disks for the DB2 installer and all the necessary software is installed on the specified hosts.

The DB2 Setup wizard can also create your DB2 instance.

The DB2 Setup wizard can also create a response file without installing the product or creating an instance. As you go through each of the GUI panels, all necessary validations are performed. Then you can update all the hostnames and shared disk names in the generated response file so that it

can be used on other hosts. You can later select any host in the host list as the installation-initiating host (IIH) to launch a response file installation.

An X server is required to display the DB2 Setup wizard.

Response file installation

A response file is a text file that contains setup and configuration values. The file is read by the DB2 setup program and the installation is performed according to the values that have been specified. A response file installation is also referred to as a silent installation.

Another advantage to response files is that they provide access to parameters that cannot be set using the DB2 Setup wizard.

If you embed the DB2 installation image in your own application, it is possible for your application to receive installation progress information in computer-readable form. This behavior is controlled by the INTERACTIVE response file keyword.

There are two ways to create a response file:

Using the DB2 Setup wizard

The DB2 Setup wizard can create a response file based on the selections you make as you proceed through the DB2 Setup wizard. Your selections are recorded in a response file that you can save to a location on your system.

One benefit of this installation method is that you can create a response file without performing an installation. This feature can be useful to capture the options required to install the DB2 product. The response file can be used at a later time to install the DB2 product according to the exact options you specified. If you plan to use the generated response file to do your DB2 pureScale Feature installation and instance setup on different hosts with different shared disks, you will need to make minor changes to that generated response file.

You can export a client or server profile with the **db2cfexp** command to save your client or server configuration, and then easily import the profile using the **db2cfimp** command. A client or server profile exported with the **db2cfexp** command can also be imported during a response file installation using the CLIENT_IMPORT_PROFILE keyword.

You should export the client or server profile after performing the installation and cataloging any data sources.

Using the sample response file

An alternative to using the DB2 Setup wizard to create a response file is to manually modify a sample response file. Sample response files are provided on the DB2 product DVD and can be found in the <image path>/db2/*platform*/samples/db2dsf.rsp, where *platform* refers to the appropriate operating system. The sample response files provide details about all the valid keywords for each product.

db2_install command

Important: The command **db2_install** is deprecated and might be removed in a future release. Use the **db2setup** command with a response file instead.

The **db2_install** command installs *all* components for the DB2 product you specify. You can then use the **db2icrt** or **db2iupdt** command to create your DB2 pureScale instance. Although the **db2_install** command installs all components for the DB2 product you specify, it does not perform user and group creation, instance creation, or configuration. This method of installation might be preferred in cases where configuration is to be done after installation. To configure your DB2 product while installing it, consider using the DB2 Setup wizard.

If you embed the DB2 installation image in your own application, it is possible for your application to receive installation progress information in computer-readable form.

This installation methods requires manual configuration after the product files are deployed.

After **db2_install** successfully installs the DB2 binary files on the installation-initiating host (IIH), you must create the DB2 pureScale instance. After the DB2 pureScale instance is created, you can add a new member or additional cluster caching facilities. For details, see the post-installation tasks.

Installing the DB2 pureScale Feature using the DB2 Setup wizard (AIX and Linux)

This topic walks you through the DB2 Setup wizard to install the IBM DB2 pureScale Feature.

Before you begin

- Ensure that your system meets installation, memory, and disk requirements.
- Complete the pre-installation tasks.
- Complete the DB2 pureScale Feature installation checklist.
- Ensure that you are logged in as root.
- Ensure you have an XServer running on the X Window client machine.
- If you are not using AIX security software to authenticate users in your environment, you must manually create the required DB2 users before you start the DB2 Setup wizard. The DB2 Setup wizard does not create the users.

About this task

The DB2 Setup wizard provides default values for most of the fields and options. You only need to input a few fields and options. For a complete list, see the installation checklist.

Procedure

1. Make the DB2 installation image available. After downloading the DB2 database product image, extract and untar the product file:
 - a. Extract the product file using the `gzip -d product.tar.gz` command, where *product* represents the name of the downloaded product.
 - b. Untar the product file using the `gnutar -xvf product.tar` command, where *product* represents the name of the downloaded product.
 - c. Change directory using the `cd ./product` command, where *product* represents the name of the downloaded product.
2. Export your display using the following command
`export DISPLAY=IP_Address:0.0`

Where *IP_Address* represents the IP address of the X Window client machine you are using to launch the installation.

3. Start the DB2 Setup wizard by running the **db2setup** command from the directory where the DB2 installation image resides:

```
./db2setup -l /tmp/db2setup.log -t /tmp/db2setup.trc
```

The **-l** and **-t** recommended parameters that enable DB2 logging and tracing. The IBM DB2 Setup Launchpad opens.

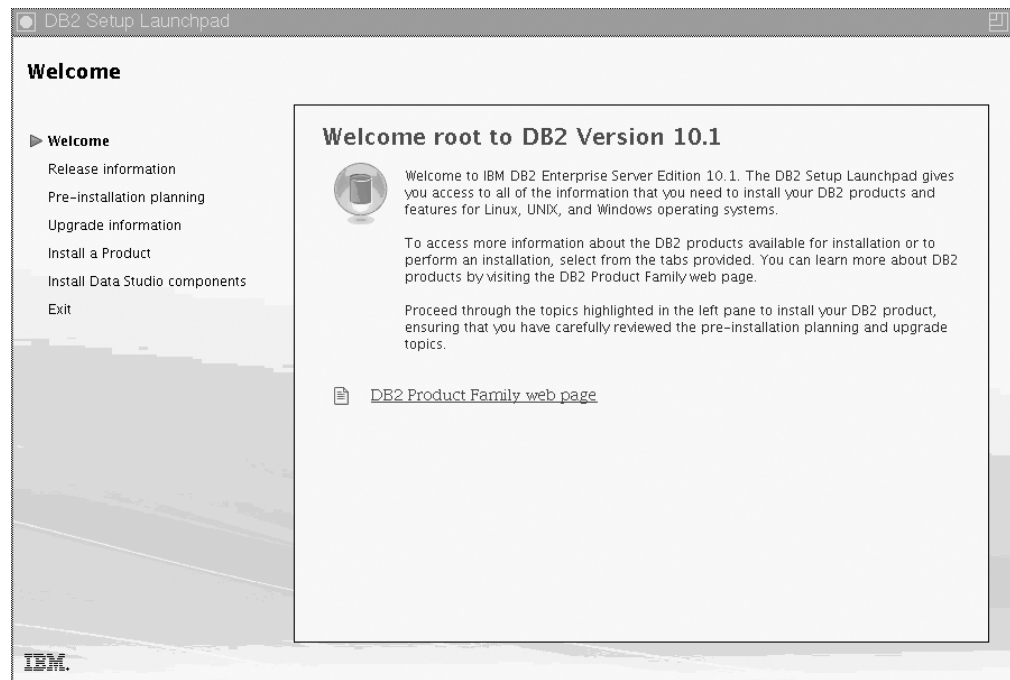


Figure 25. DB2 Setup Launchpad

From this window, you can view the DB2 pureScale Feature release information, upgrade options, architectural overview, and pre-installation planning, or you can proceed directly to the installation.

4. Click **Install a Product** and the Install a Product window displays the products available for installation.




Figure 26. Install a Product window

Click **Install New** under DB2 Enterprise Server Edition Version 10.1 with the IBM DB2 pureScale Feature. Installation help can guide you through each of the installation panels. To invoke the installation help, click **Help** or press **F1**. You can click **Cancel** at any time to end the installation.

5. Welcome Panel.




Figure 27. DB2 Setup wizard Welcome Panel

Click **View Features** to view the features that will be installed. The View Features dialog box opens and displays the features to be installed. Click **Close** to close the View Features dialog box.

Click **Next** to continue.

6. Software License Agreement Panel.

After reviewing the software licensing agreement, select **Accept** to accept the




Figure 28. Software License Agreement Panel

terms. Click **Next**.

7. Select installation, response file creation, or both Panel




Figure 29. Select installation, response file creation, or both Panel

You can use the DB2 Setup wizard to:

- Install the product. When you are finished stepping through the DB2 Setup wizard, your DB2 database product is installed on your system. No response file is saved.
- Create and save the installation settings in a response file, but not install the product. You can use this response file to install this product later. The DB2 database product is not installed. The DB2 Setup wizard generates the response file only based on your selections.
- Install the product and create and save a response file that you can use as a record of your choices.

If you choose to create a response file, you can specify the name and location of the response file in the **Response file name** field. This same response file can be updated to use on another set of hosts as well. Click **Next**.

8. Select installation directory Panel



Figure 30. Select installation directory Panel

Specify the directory path where you want to install your DB2 database product. If you are installing the DB2 pureScale Feature, the DB2 installer installs the DB2 copy to the same location on each host and this specified location must be valid on all hosts.

The DB2 database resides on the shared disk and the DB2 code is installed on all DB2members.

Click **Next**.

9. Select the language to install Panel




Figure 31. Select language to install Panel

Specify the directory path where you want to install your DB2 languages. If you are installing the DB2 pureScale Feature, the DB2 installer installs the DB2 copy to the same location on each host and this specified location must be valid on all hosts.

Click **Next**.

10. Specify the location of the DB2 Information Center Panel

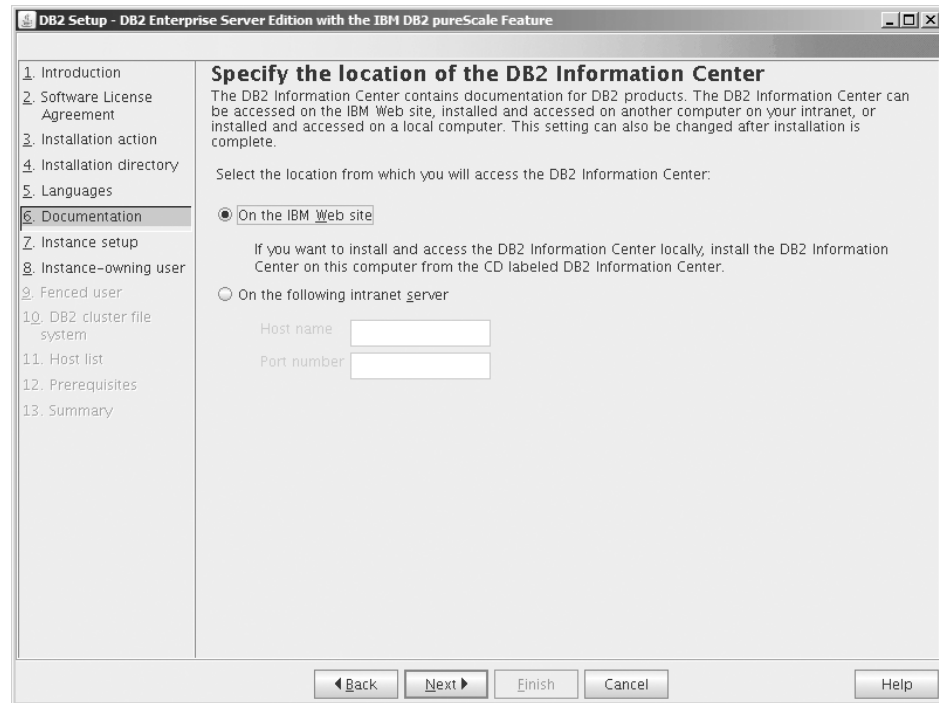


Figure 32. Specify the location of the DB2 Information Center Panel

Specify the directory path of your DB2 Information Center. You can select the IBM website or specify the intranet server with the host name and port number.

Click **Next**.

11. Set up a DB2 instance Panel.

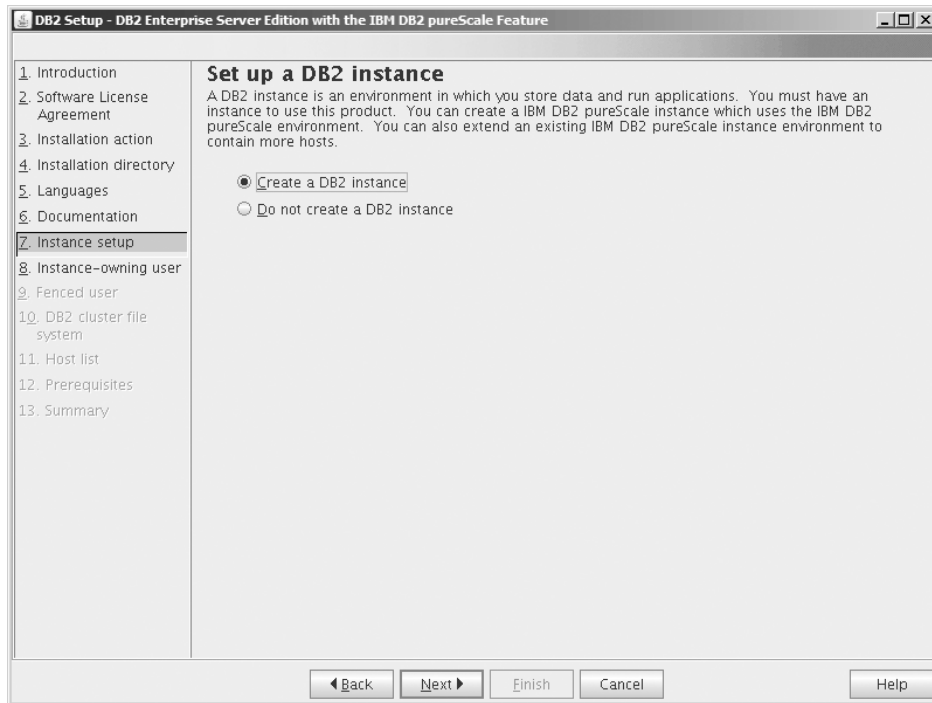


Figure 33. Set up a DB2 instance Panel

Use this panel to create a DB2 pureScale instance. If you would like to add the current host to an existing DB2 pureScale instance, you should start the instance extension (**db2isetup** or **db2iupdt -add**) from a host that already belongs to the instance.

- To create a DB2 instance, select **Create a DB2 instance**.
- If you do not want the DB2 Setup wizard to create a DB2 instance, select **Do not create a DB2 instance**. The DB2 installer still performs the installation, but, you can create an instance at a later point by running the **db2icrt** or **db2isetup** command.

Click **Next**.

12. Set user information for the DB2 instance owner Panel

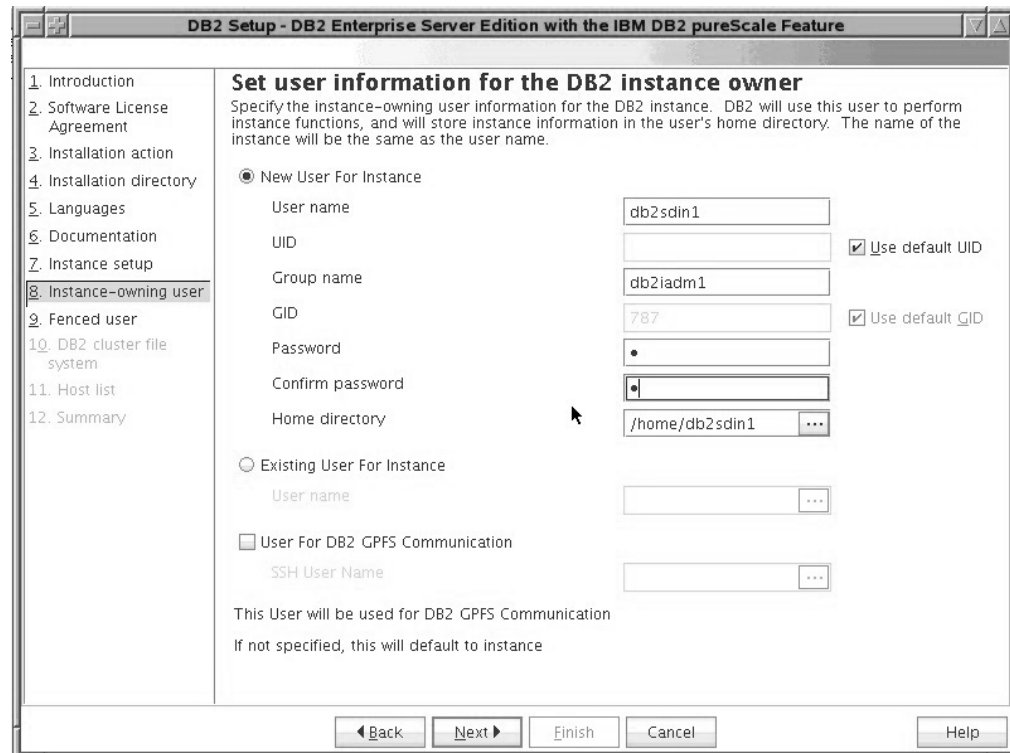


Figure 34. Set user information for the DB2 instance owner Panel

Specify the user information that the DB2 installer uses to perform instance functions and store instance information in the user's home directory. The new user and group is created on all hosts. The name of the instance is the same as the user name. You can also specify an existing user. If you select "Existing user", the user name must comply with the conditions outlined in the "Required users for a DB2 pureScale Feature installation" topic. The SSH User Name is required to use a secure shell (SSH) network protocol between hosts. If this name is not specified, the instance owner user ID is used. Click **Next**.

13. Set user information for the fenced user.

DB2 Setup - DB2 Enterprise Server Edition with the IBM DB2 pureScale Feature

1. Introduction
2. Software License Agreement
3. Installation action
4. Installation directory
5. Languages
6. Documentation
7. Instance setup
8. Instance-owning user
9. Fenced user
10. DB2 cluster file system
11. Host list
12. Prerequisites
13. Summary

Set user information for the fenced user

Specify the required information for the fenced user. Fenced user defined functions (UDFs) and stored procedures will execute under this user and group.

☒ **New user**

User name:

UID:

Group name:

GID:

Password:

Confirm password:

Home directory:

☒ Use default UID

☒ Use default GID

☐ **Existing user**

User name:

Back Next Finish Cancel Help

Figure 35. Set user information for the fenced user Panel

Create a new user or use an existing user:

- To create a new fenced user, select **New user** and provide information about the fenced user in the fields. In a DB2 pureScale environment the user account is created on all the hosts.
- To use an existing fenced user, select **Existing user**. You can type a valid User name or click [...] to select an existing user name. If you select "Existing user", the user name you enter must exist on all hosts with the same UID and GID.

Click **Next**.

14. Setting up a DB2 Cluster File System Panel

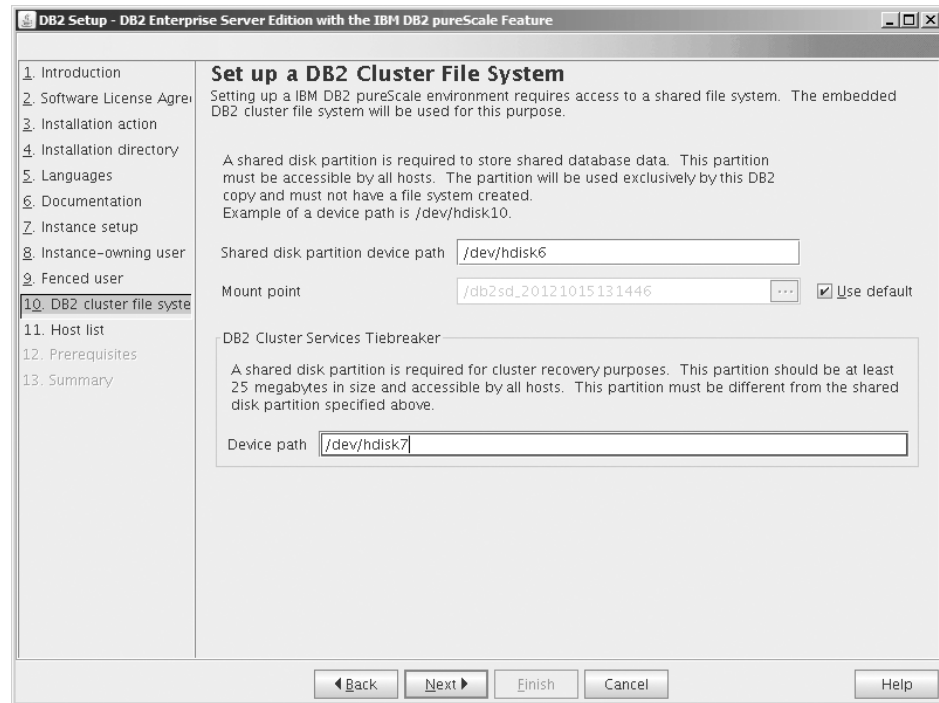


Figure 36. Setting up a DB2 Cluster File System

Use this panel to specify both of the shared disk partitions required for the DB2 pureScale Feature installation:

- One partition where the database data resides (Shared disk partition device).
- One partition for cluster recovery purposes (DB2 Cluster Services tiebreaker).

The partitions specified cannot be the same. They should be raw disks and not be used by anything else. Click **Next**.

15. Host List Panel




Figure 37. Host List Panel

Add the hosts that are part of the DB2 pureScale environment. Hosts are validated as they are added to the list. The available disk space for the installation files is checked on every host as part of the validation process. The installation-initiating host (IIH) is automatically loaded into the host list. To add a remote host, click **Add**.

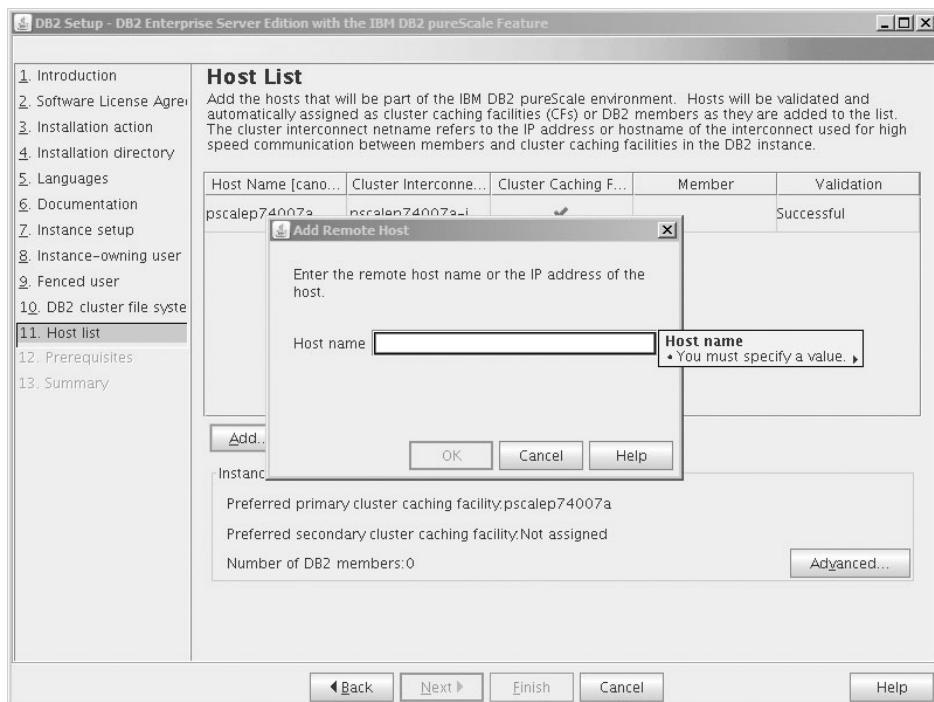


Figure 38. Add Remote Host window

Enter the remote host name and click **OK**. If you receive an error during host remote validation, you can fix the problem on the Host List panel or return to any of the previous panels of DB2 Setup wizard to make changes to the installation settings. You can also click **Advanced** on the Host List page to change the DB2 pureScale instance settings. In addition, regardless of the status of the hosts, a secure shell check is done at this point to confirm the passwordless SSH availability between the hosts. Ensure that root SSH access is installed, setup, and configured on the host you want to add.

Error information is available by clicking the button displayed in the "Errors" column. If you cannot resolve a host error, you can remove the host from the list entirely with the **Remove** button. You can then later fix the issue and add the host to the instance after the initial deployment. You can also specify custom settings for the DB2 pureScale instance. See the next step for more details.

16. Advanced Instance Settings dialog. After adding all the hosts you want to participate in the DB2 pureScale environment, you can make changes to the default instance settings by clicking on **Advanced**. Optionally, you can make changes to the default instance settings using the options in the **Instance Setting** box.

The following tabs can be displayed:

Cluster caching facility

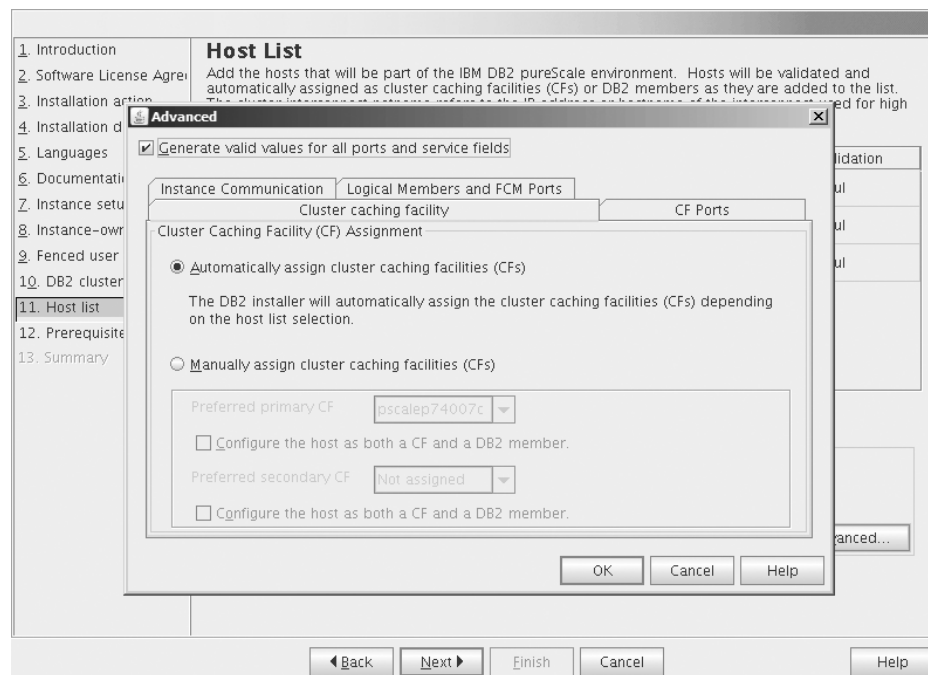


Figure 39. Advanced Instance Settings - Cluster caching facility tab

You can manually select the hosts to be used as CFs, or allow the installer to automatically assign them. The automatic CF assignment option is selected by default. At least one CF is required to run a DB2 pureScale instance.

If you choose to manually assign the CF, you need to specify which CF acts as the preferred primary CF. You can also optionally designate

which CF acts as the preferred secondary CF. These designations indicate which cluster caching facility the cluster manager attempts to start in the primary CF role.

You can also select the option of specifying that a host can act as both a CF and a member.

CF Ports

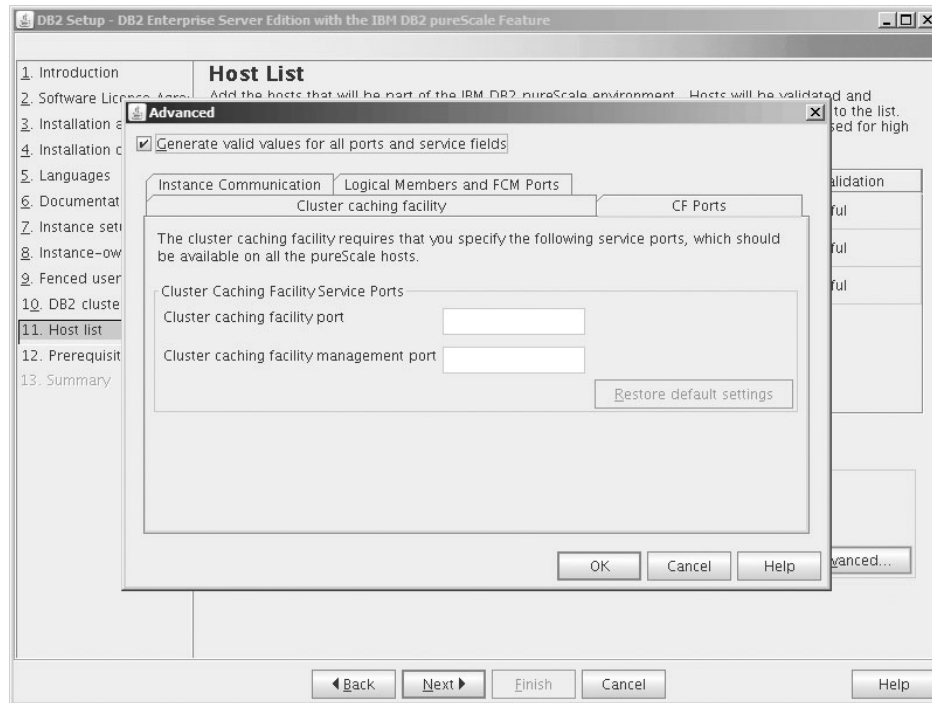


Figure 40. Advanced Instance Settings - CF Ports tab

The CF requires that you specify service ports. The same ports need to be available on all hosts that are running cluster caching facilities. The port numbers can be in the range of 1024 to 65535.

Instance Communication

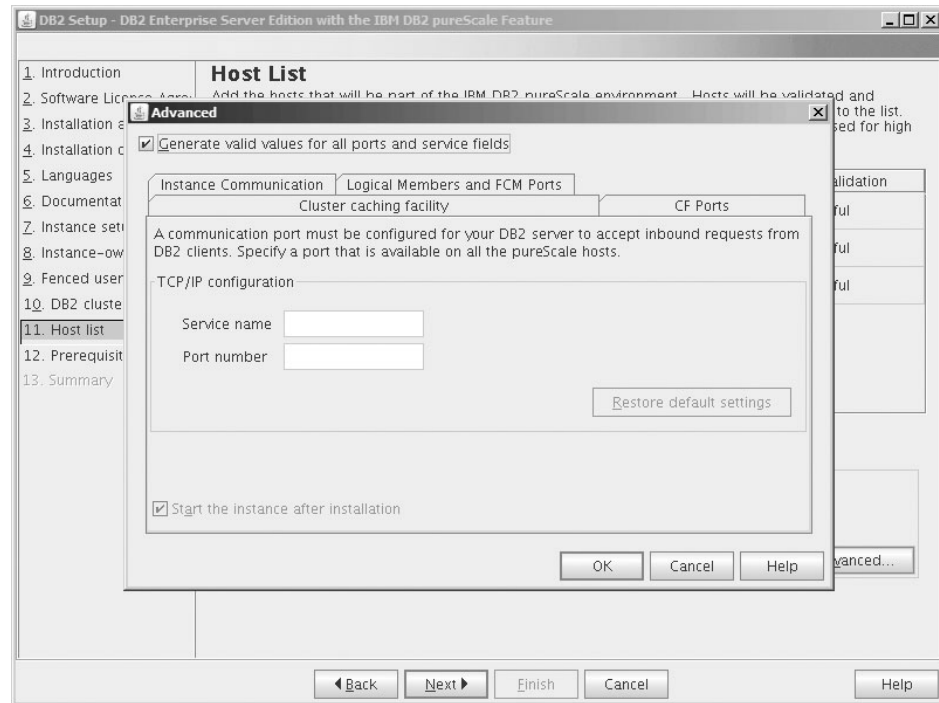


Figure 41. Advanced Instance Settings - Instance Communication tab

The TCP/IP port number is used by a DB2 instance to listen for incoming DB2 connections. The port number value must be in the range 1024 to 65535 and must be available on all hosts.

Logical Members and FCM Ports

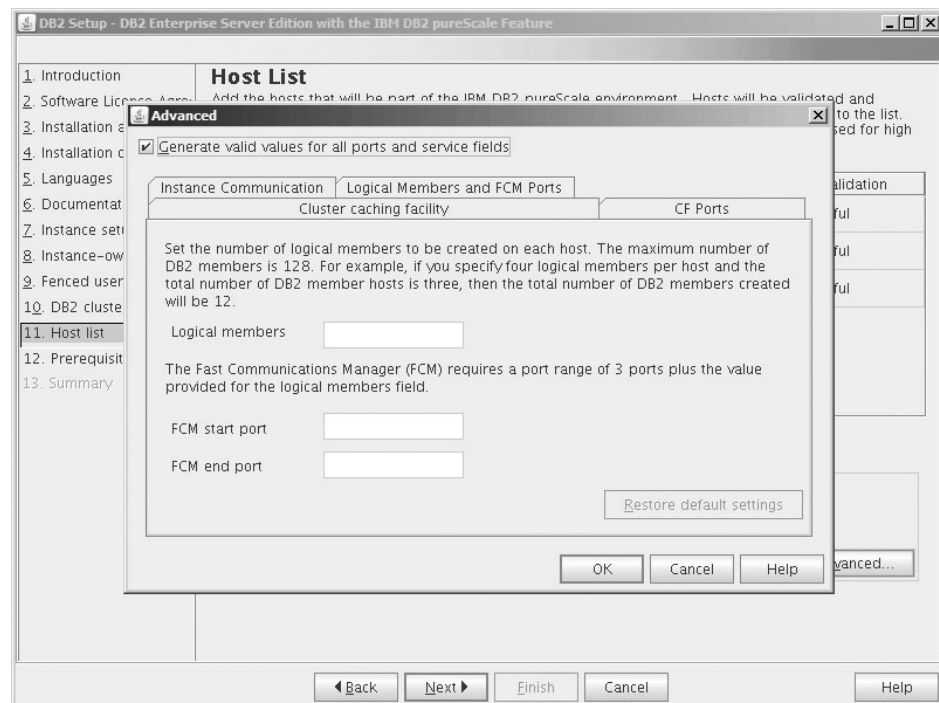


Figure 42. Advanced Instance Settings - Logical members and FCM Ports tab

You can set the number of logical members to be created on each host. The maximum number of members is 128. The same logical number is applied to all hosts. Ensure that the port numbers used by the DB2 pureScale instance are free on all hosts. The Fast Communications Manager (FCM) requires a port range of the three mandatory ports plus the value provided for the logical members field. This port range can designate up to 130 hosts (128 member + 2 CFs). The default FCM start port is 60000 and must be in the range of 1024 to 65535.

Click **OK** to close the Advanced dialog. The DB2 installer revalidates the added hosts if you have updated any of the default instance settings. Click **Next** on the Host List Panel.

17. **DB2 Prerequisite Panel** Use this panel to determine if your system satisfies the DB2 pureScale Feature prerequisites (such as operating system level, Linux distribution, AIX Technology Level, and other DB2 pureScale specific requirements.) The status field indicates whether or not the prerequisite are met. Click **Next**.
18. **Summary Panel**
You have made all the choices necessary to complete the DB2 installation or to

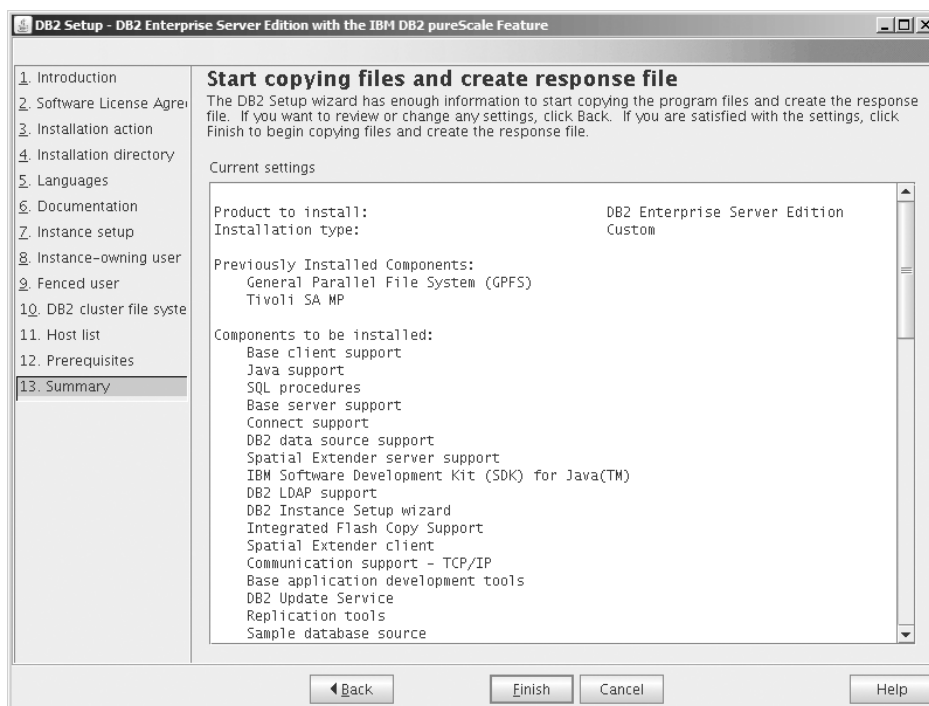


Figure 43. Summary Panel

create a response file. Review all the features listed in this window before clicking **Finish**. To change any of these settings, click **Back** to go back to the previous pages in the DB2 Setup wizard and change your choices. Clicking **Finish** launches the installation process across all specified hosts.

19. **Setup Complete**

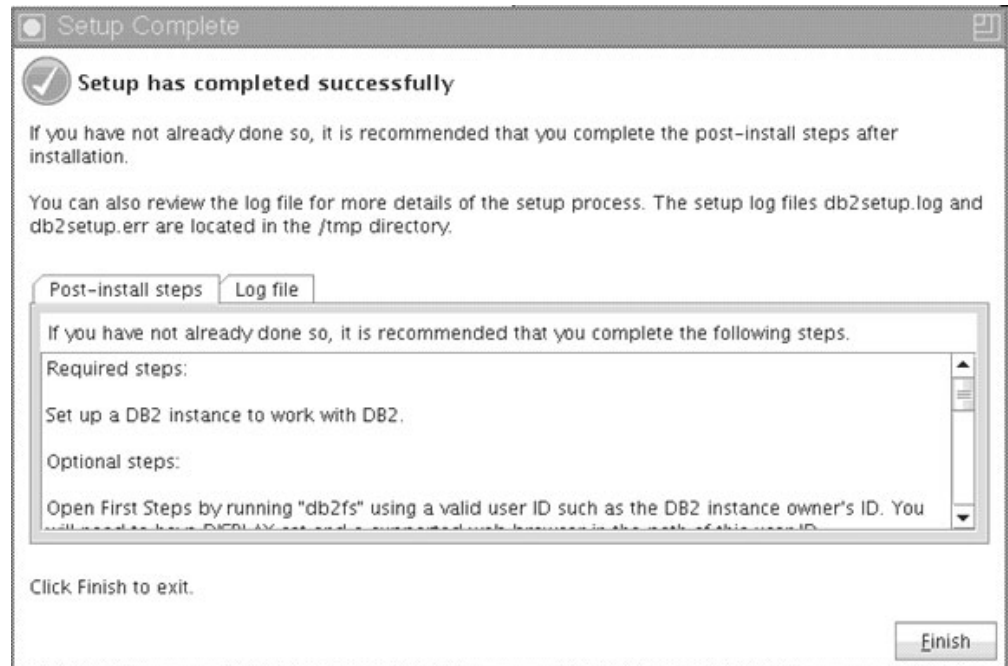


Figure 44. Setup Complete Panel

Once the installation is complete, you can review the post-installation steps, the log file for more details on what was setup, and the status of each remote host. Click **Finish** to exit.

Results

The DB2 pureScale Feature installation image is copied to the DB2DIR/sd directory on each host in the DB2 pureScale instance. If you need to scale the DB2 pureScale instance, you can use any host within the DB2 pureScale instance from which to launch the operation.

The default names and location for the installation log files are:

- The DB2 setup log file. This file captures all DB2 installation information including errors. The DB2 setup log file name is db2setup.log.
- The DB2 error log file. This file captures any error output that is returned by Java (for example, exceptions and trap information). The DB2 error log file name is db2setup.err.

By default, these log files are located in the /tmp directory. You may have overridden this default when you issue the **db2setup** command. The DB2 installer saves a copy of the DB2 setup log file in the DB2DIR/install/logs/ directory, and renames it db2install.history. If the name already exists, then the DB2 installer renames it db2install.history.xxxx, where xxxx is 0000-9999, depending on the number of installations you have on that machine.

The DB2 installer also produces logs detailing the actions on each host, including the ITH, under the /tmp/db2log.xxxx directory on the local machine (where xxxxxx represents the process id) and /tmp/db2log directory on each remote host.

When you want to add another DB2 member or CF, you can run either the **db2isetup** or **db2iupdt -add** command from any of the existing DB2 member or

CF. These commands performs the DB2 installation and instance configuration on the target hosts.

What to do next

Review and perform any necessary post-installation tasks.

Installing a DB2 database product using a response file (Linux and UNIX)

These instructions guide you through the process of installing a DB2 database product on a Linux or UNIX operating system using a response file.

Before you begin

Before you begin the installation:

- Ensure that your system meets installation, memory, and disk requirements.
- Complete the pre-installation tasks.
- Complete the IBM DB2 pureScale Feature installation checklist.
- Ensure that you are logged in as root.
- For root installations, log on with the user ID that has root privileges. For non-root installations, log on with the user ID that is to own the DB2 installation.
- Your system meets all of the memory, hardware, and software requirements to install your DB2 database product.
- You can install a DB2 database product using a response file with either root or non-root authority.
- All DB2 processes associated with the copy of DB2 you are working with are stopped.
- Ensure the response file keyword LIC_AGREEMENT is set to ACCEPT.

Procedure

To install the DB2 pureScale Feature using a response file:

1. Enter the **db2setup** command as follows:

```
./db2setup -r responsefile_directory/response_file
```

where:

- *responsefile_directory* represents the directory where the response file is located; and
 - *response_file* represents the name of the response file.
2. Check the messages in the log file when the installation finishes. The installation logs are located, by default, in the */tmp* directory:
 - *db2setup.log* (*db2setup_username.log* for non-root installations, where *username* is the non-root user ID under which the installation was performed)
 - *db2setup.err* (*db2setup_username.err* for non-root installations, where *username* is the non-root user ID under which the installation was performed)

You can specify the location of the log file. An installation history log *db2install.history* is located in *DB2DIR/install/logs* where *DB2DIR* specifies the path where the DB2 database product was installed. The default installation directory is:

- For AIX, HP-UX or Solaris /opt/IBM/db2/V10.1
- For Linux/opt/ibm/db2/V10.1

If more than one DB2 database product is installed in the same location, you will see `db2install.history.xxxx`, where `xxxx` are digits, starting at 0001 and increasing depending on how many DB2 copies you have installed.

Example

The following example installs the DB2 pureScale Feature using response file `db2dsf.rsp`:

```
db2setup -r /mydir/db2dsf.rsp -l /tmp/db2setup.log -t /tmp/db2setup.trc
```

where `mydir` represent the directory where you saved the `db2dsf.rsp` file.

What to do next

If you want your DB2 database product to have access to DB2 documentation either on your local computer or on another computer on your network, then you must install the *DB2 Information Center*. The *DB2 Information Center* contains documentation for the DB2 database and related products.

Review and perform any necessary post-installation tasks.

Installing DB2 pureScale Feature using the `db2_install` command

This topic guides you through the process of manually installing IBM DB2 pureScale Feature using the `db2_install` command.

Before you begin

Important: The command `db2_install` is deprecated and might be removed in a future release. Use the `db2setup` command with a response file instead.

- Ensure that your system meets installation, memory, and disk requirements.
- Complete the pre-installation tasks.
- Complete the DB2 pureScale Feature installation checklist.
- Ensure that you are logged in as root.

About this task

The manual installation process is outlined as follows:

1. Run the `db2_install` command to install the DB2 database product on the installation-initiating host (IIH)
2. Run the `db2icrt` command on the IIH to create one member and one cluster caching facility. The `db2icrt` command will copy the DB2 database product to the second host and run the install process before creating the member and cluster caching facility.
3. Run the `db2iupdt` to add members and cluster caching facilities. The `db2iupdt` command will copy the DB2 database product to the new host and install it before creating the member or the CF.

Procedure

1. Make the DB2 installation image available. After downloading the DB2 database product image, decompress and untar the product file:
 - a. Decompress and untar the product file using the `gunzip -c product.tar.gz | tar -xf -` command, where *product* represents the name of the downloaded product.
 - b. Change directory using the `cd ./product` command, where *product* represents the name of the downloaded product.
2. On the I/H, run the **db2_install** command to install the DB2 pureScale Feature. This command will install the DB2 database product, Tivoli SA MP, and GPFS.

Note:

For example:

```
db2_install -b DB2DIR -p ese_dsf -t /tmp/db2_install.trc  
-l /tmp/db2_install.log
```

where *DB2DIR* represents the location where you want to install the DB2 pureScale Feature image. This path has to be the same on all the hosts. If the *DB2DIR* exists, it must be empty before the installation. The **-t** and **-l** options are recommended so that trace and log files are generated.

3. Verify that your DB2 database product is installed. Run the `/usr/local/bin/db2ls` command.
4. Verify that Tivoli SA MP and GPFS are installed on the system.
 - a. To verify the installation of Tivoli SA MP, run the `rpm -aq | grep "sam\."` | `more` command. The output of this command should be similar to the following example:

```
hostinst5:/ # rpm -aq | grep sam | grep -v "sam\."  
sam-3.2.2.1-11349  
sam.adapter-3.2.2.1-11349
```
 - b. To verify the installation of GPFS, run the `rpm -aq | grep gpfs` command. The output of this command should be similar to the following:

```
hostinst1:/opt/ibm/db2/V10.1/sd/db2/linuxamd64/gpfs # rpm -aq | grep gpfs  
gpfs.msg.en_US-3.5.0.4  
gpfs.base-3.5.0.4  
gpfs.docs-3.5.0.4  
gpfs.gpl-3.5.0.4
```
5. Run the following command as root to set up the initial DB2 pureScale instance specifying a host as a DB2 member and another host as a cluster caching facility. You can find the **db2icrt** command in the *DB2DIR/instance* directory. You can only specify one DB2 member and one cluster caching facility during the initial setup of the DB2 pureScale instance. You can add additional DB2 members and cluster caching facilities using the **db2iupdt** command.

Note:

- The netname represents the cluster interconnect netname which corresponds to the communication adapter port used for high speed communication between DB2 members and cluster caching facilities in the DB2 instance.
- For the host name, you need to use the same host name format as the return value of the **hostname** command

```
db2icrt -d -m member_hostname:member_netname1  
-cf CF_hostname:CF_netname2  
-instance_shared_dev disk1_name -tbdev IP_address  
-u fencedID instname
```

For example,

```
db2icrt -d -m h1.domain.com:h1.domain.com-ib0
-cf h2.domain.com:h2.domain.com-ib0
-instance_shared_dev /dev/hdisk1 -tbdev 129.42.38.1
-u db2fenc1 db2sdin1
```

will create DB2 pureScale instance db2sdin1 with a DB2 member named h1.domain.com, a netname of h1.domain.com-ib0, and cluster caching facility h2.domain.com with a netname of h2.domain.com-ib0. It will also specify the shared file system device path as /dev/hdisk1, and set the tiebreaker disk as /dev/hdisk2.

What to do next

You can add other hosts to your DB2 pureScale instance using the **db2iupdt -add** command.

Setting up a Network Time Protocol server

This topic outlines how to setup a Network Time Protocol (NTP) server to ensure that all servers that the instance runs on have similar time settings.

By default, the NTP server is setup by the DB2 installer. During instance creation and update, the NTP is automatically setup on the primary CF. This setup is a best practice that will help with problem diagnosis and informal server monitoring.

Before you begin

In a DB2 pureScale environment, to synchronize operations and facilitate time sensitive operations, the system clocks of all members must be synchronized.

For this reason, NTP must be setup on all of your members. The DB2 installer installs and sets up the NTP server and clients. By default, the NTP server is the primary CF. You can use the steps in this topic to change that default setting and assign another host to act as the NTP server.

The DB2 database manager checks whether NTP is correctly set up on each member and verifies the presence of the NTP daemon. The DB2 database manager also periodically compares time stamps between members to check that they differ by less than the value set by the MAX_TIME_DIFF database manager configuration parameter.

In a DB2 pureScale environment, the value of MAX_TIME_DIFF represents the maximum number of minutes allowed between the system clocks of two members; the default is 1 minute. For each member, if NTP is not correctly set up or the NTP daemon is not present, or if time stamps between members differ by more than the value of MAX_TIME_DIFF, warnings are written to the **db2diag** log files.

About this task

If you do not have access to an NTP time server, use the following steps to configure one of the hosts in the cluster or another host accessible by members as an NTP server. The system clock on this host will be used as the clock that is synchronized. This host must remain online to continue synchronizing the clocks on the hosts.

Procedure

1. Choose the host that you will setup as an NTP server.
2. Setup the NTP daemon to synchronize to its own system clock. Edit `/etc/ntp.conf` and add the following lines:

```
server 127.127.1.0 prefer # the logical IP address for the NTP server to sync
to its own system clock
fudge 127.127.1.0
driftfile /etc/ntp.drift
tracefile /etc/ntp.trace
```
3. Configure the NTP server to start at system restart.
 - On AIX operating systems, edit the file `/etc/rc.tcpip` and uncomment the following entry:

```
start /usr/sbin/xntpd -x "$src_running"
```
 - On Linux operating systems, run the **chkconfig ntp on** command.
4. Start the NTP server:
 - On AIX operating systems, run the `startsrc -s xntpd` command.
 - On Linux operating systems, run the **service ntp start** command.

Configuring hosts as Network Time Protocol clients

This topic outlines how to configure a Network Time Protocol (NTP) server to ensure that all servers that the instance runs on have similar time settings.

About this task

By default, the NTP server is setup and configuration is handled by the DB2 installer.

Procedure

To configure hosts as NTP clients:

1. Specify the NTP server to synchronize to. Open the `/etc/ntp.conf` file and add the following lines:

```
server ntp_server_hostname
driftfile /etc/ntp.drift
tracefile /etc/ntp.trace
```

where `ntp_server_hostname` represents the hostname or IP address of an NTP server. If you do not have access to an existing NTP server, refer to the “Setting up an NTP Server” topic.
2. Configure the NTP server to start at system restart:
 - On AIX operating systems, edit the `/etc/rc.tcpip` file and uncomment the following line:

```
start /usr/sbin/xntpd -x "$src_running"
```
 - On Linux operating systems, run the **chkconfig ntp on** command.
3. Start the NTP server:
 - On AIX operating systems, run the **startsrc -s xntpd** command.
 - On Linux operating systems, run the **service ntp start** command.
4. Verify that the daemon is synchronized. If the daemon is significantly not synchronized with the system clock, it can take over 10 minutes for it to become synchronized. To ensure the daemon is synchronized before continuing you can check the “sys stratum” field:

- On AIX operating systems, use the output of the `lssrc -ls xntpd` command.
- On Linux operating systems, use the output of the `ntptrace` command.

The `sys stratum` field should be less than 16 to continue to the next step.

5. Finish the synchronization process by running the `ntpdate -d ntp_server_hostname` command.

Chapter 18. Taking the first steps after installing the DB2 pureScale Feature

After you install a DB2 pureScale Feature, there are additional necessary and optional tasks available.

Verifying your DB2 pureScale installation and setup

Use the steps in this task to verify your installation and setup of the IBM DB2 pureScale Feature.

About this task

The following steps outline a procedure for you to verify your DB2 pureScale installation and setup. Perform these steps after you have created your DB2 pureScale instance and added all of the members and CFs you want in your DB2 pureScale environment.

Procedure

1. Check the installation logs for any errors.
2. As the instance user, run the **db2instance - list** command. This command returns a list of the members and cluster caching facilities (CFs) in your DB2 pureScale environment. If you manually installed the DB2 pureScale Feature, the status of the members and CFs in your DB2 pureScale environment will be in the STOPPED state. You can, run the **db2start** command to start the DB2 pureScale instance and rerun the **db2instance - list** command.
3. Verify that all the members and CFs have been started (STATE = STARTED). One CF should be PRIMARY. The other CF should be in either PEER state or CATCHUP state.

Note: When you view the status of the non-primary CF using the command **db2instance -list**, it will be in **CATCHUP** state until a connection is made to the database. Once the first connection is made, the process of copying data from the primary CF begins.

4. Run the **db2samp1** command to create a sample database.
5. Issue the following command to connect to the SAMPLE database:

```
db2start
db2 connect to SAMPLE
```

6. Issue the following statement:

```
db2 "select * from EMP"
```

If this command returns records, then everything is working as it should.

Manually allowing the DB2 installer to take over a user-managed GPFS cluster

Use this task to pass the ownership of your user managed GPFS cluster to the IBM DB2 pureScale Feature.

Before you begin

- You need root access on each host to perform these steps.
- The DB2 pureScale Feature must be installed on all hosts in the GPFS cluster

About this task

Only GPFS clusters created by DB2 data server products are registered in the DB2 global registry. Performing this task will take your existing user managed GPFS cluster and treated it as if it had been created by a DB2 data server product. The DB2 installer will then perform the clustered file system tasks during DB2 pureScale instance operations and update the GPFS binaries when running the **installFixPack** command.

Procedure

To take over a user-managed GPFS cluster:

1. Log on as root on any machine in your cluster.
2. Run the **db2cluster_prepare** command with the following parameters:
`db2cluster_prepare -cfs_takeover`
3. Check the error return code using the **echo \$?** command. If there are errors, resolve the errors and rerun the command as specified in Step 2.
4. To verify that you've properly added the record, run the following command:
`db2dir/bin/db2greg -dump`

The following sample output should be returned:

```
V,GPFS_CLUSTER,NAME,db2cluster_20111117140104.torolab.ibm.com,-,DB2_MANAGED
```

What to do next

You can now proceed with the installation of the DB2 pureScale Feature.

Enabling SCSI-3 PR for DB2 pureScale Feature (Linux)

You can enable SCSI-3 PR for the IBM DB2 pureScale Feature which will provide faster fail-over support.

Before you begin

This task should be performed after you have installed the DB2 pureScale Feature and created your instance. If you already have GPFS installed and have a GPFS cluster running , for example, if you have DB2 pureScale Feature Fix Pack 2 installed on your SLES hosts, then after moving to DB2 pureScale Feature Fix Pack 3, you can perform this task at any time before using the DB2 pureScale environment.

You need instance user and root user authority.

About this task

Not all disks support SCSI-3 PR on Linux.

Procedure

1. Log in as the instance user.
2. Stop the database manager by running the **db2stop** command.
3. Log in as root user.
4. Stop GPFS on all hosts by running the **db2cluster -cfs -stop -all** command.

5. Find the filesystems as the instance user:

```
coralm201:/home/madhusuk >db2cluster -cfs -list -filesystem
```

FILE SYSTEM NAME	MOUNT_POINT
fs	/fs

6. Find the disk info for that filesystem as the instance user:

```
coralm201:/home/madhusuk >db2cluster -cfs -list -filesystem fs -disk
```

PATH ON LOCAL HOST	OTHER KNOWN PATHS
(*) /dev/sdc	

7. Find the information for that disk by running the **tsprinqiry** command as root:

```
coralm201:# /usr/lpp/mmfs/bin/tsprinqiry sdc
```

```
IBM      :VirtualDisk      :0730
```

8. Create the `/var/mmfs/etc/prcapdevices` file on each host in the DB2 pureScale environment and edit the file to add your disk. For example, using the output from **tsprinqiry** command:

```
IBM:VirtualDisk:0730
```

9. Tell GPFS to use SCSI-3 PR by running the **/usr/lpp/mmfs/bin/mmchconfig usePersistentReserve=yes** command.

10. As the root user, start GPFS on all nodes by running the **db2cluster -cfs -start -all** command.

11. Once the GPFS filesystems are mounted on all hosts, check that SCSI-3 PR has been properly enabled by running the **/usr/lpp/mmfs/bin/mm1nsd -X** command on all machines and checking that "pr=yes" appears in the Remarks column. For example:

```
coralm201:/usr/lpp/mmfs/bin # ./mm1nsd -X
```

Disk name	NSD volume ID	Device	Devtype	Node name	Remarks
gpfs12nsd	091A5CBD4C6B8076	/dev/sdc	generic	coralm201.torolab.ibm.com	pr=yes

12. As the instance user, start the database manager by running the **db2start** command.

Results

You can now use the DB2 pureScale Feature with SCSI-3 PR.

Registering a DB2 product or feature license key using the **db2licm** command

Use the **db2licm** command to apply the license entitlement certificate (also referred to as registering a license key.)

Before you begin

The IBM DB2 pureScale Feature installation image includes a copy of the DB2 Enterprise Server Edition code, however, the installation image does not come with a license entitlement for DB2 Enterprise Server Edition. You must have purchased a separate entitlement for DB2 Enterprise Server Edition sufficient to cover all users or processors in your cluster, and register the DB2 Enterprise Server Edition and feature licenses. (The DB2 pureScale Feature license is installed automatically by the DB2 installer.)

To complete this task, you must have the DB2 Enterprise Server Edition license or feature license. The license key can be found on the Activation CD, which is included in your product package or can be downloaded from the Passport Advantage website (<http://www.ibm.com/software/lotus/passportadvantage/>.)

About this task

In a DB2 pureScale environment, before adding new members, register all licenses on the installation-initiating host (IIH.) In this case, if you initiate installation onto more members from the IIH the DB2 installer automatically copies all license certificate files onto the new members and installs them. However, if you install new members from their own machine, you must manually install the license certificate files on each member using the **db2licm** command.

If you install additional licenses after the new member has been added, you must register the license certificate file on each member manually.

Procedure

Register a DB2 license key by entering the following command:

```
/opt/IBM/db2/V10.1/adm/db2licm -a filename
```

where *filename* is the full path name and file name for the license file that corresponds to the product or feature.

Creating a DB2 pureScale instance

This topic details the steps necessary to create a DB2 pureScale instance in a DB2 pureScale environment.

Before you begin

- Login as root to a host where the DB2 pureScale Feature was installed.
- Ensure that your network topology is set up according to the requirements in the Chapter 12, “Network topology configuration considerations for DB2 pureScale environments,” on page 145 topic.
- Ensure all hosts in the DB2 pureScale environment are at the latest DB2 fix pack level.

About this task

If you installed DB2 pureScale Feature by using either the DB2 setup wizard or a response file, by default a DB2 pureScale instance is created. However, if you choose to not create the instance during installation, or, if you installed DB2 pureScale Feature by using the **db2_install** command, you must manually create the DB2 pureScale instance.

Important: The command **db2_install** is deprecated and might be removed in a future release. To install with the command line, run the **db2setup** command with a response file to specify the installation options. To specify the response file, use the **-r** option and the response file path. The **db2setup** command has the benefit of being able to create a DB2 pureScale instance, and set up the cluster caching facilities and DB2 members.

Use this topic to create a DB2 pureScale instance after you install the DB2 pureScale Feature.

You can use the **db2icrt** command to create a DB2 pureScale instance with one cluster caching facility (CF) and one member. The **db2icrt** command must be issued on either the target CF host or the target member host. The host where the command is issued is referred to as installation-initiating host (IIH).

You can also use the **db2isetaup** command to create an instance. The **db2isetaup** command can also be used to create an instance with multiple hosts, whereas the **db2icrt** command can create an instance only on two hosts. Regardless of the command used to create the instance, additional hosts can be added by running the **db2iupdt** command with the **-add** option.

Restrictions

When you create a DB2 pureScale instance by using the **db2icrt** command, you must specify one member and one CF.

Procedure

To create a DB2 pureScale instance by using the **db2isetaup** or **db2icrt** command:

- To create a DB2 pureScale instance by using the **db2isetaup** command:
 1. Run the **db2isetaup** command to launch the DB2 Instance Setup wizard.
 2. Proceed to the “Set up a DB2 instance” panel. Select the Create a DB2 instance option and click **Next**.
 3. Add the instance owner and the fenced user in the “Set user information” panels and click **Next**.
 4. On the “Set up a DB2 Cluster File System” panel, specify the shared disk partition device path and the tiebreaker device path.
 5. Proceed to the Host List panel. Add the hosts that you want in the DB2 pureScale instance configuration. Hosts are validated as they are added to the list. After you add the hosts, click **Next**.
 6. Click **Finish** to perform the action.

- To create a DB2 pureScale instance by using the **db2icrt** command:

Run the **db2icrt** command to create the DB2 pureScale instance. The command syntax is as follows:

```
db2icrt -m <MemberHostName> -mnet <Netname1>  
        -cf <CFHostName> -cfnet <CFNetname>  
        -instance_shared_dev <Shared_Device_Path_for_Instance>  
        -tbdev <Raw_device_for_tiebreaker>  
        -u <fencedID>  
        <instanceID>
```

This syntax diagram is abbreviated for simplicity, see the **db2icrt** command reference for the full syntax and parameter descriptions. One of the host names specified must be the name of the host that you are running the **db2icrt** command from.

Note: If you have not set your communication protocols correctly, you might receive a SQL5043N error when issuing the **db2icrt** command. Ensure that you have correctly set the **DB2COMM** environment variable. Use the **db2set -all** command to view your environment settings. You can set the **DB2COMM** environment variable by running the following command: **db2set DB2COMM=tcPIP**.

Example

For example, use the **db2icrt** command to create a DB2 pureScale instance called *db2sdin1* on the installation-initiating host, *test1*, and *test2* (cluster caching facility) that use the shared disk */dev/hdisk1* (identifier on *test1*) as the tiebreaker and an existing file system */db2instance* as the instance shared disk:

```
DB2DIR/instance/db2icrt
-m test1 -mnet test1-ib0
-cf test2 -cfnet test2-ib0,test2-ib1,test2-ib2,test2-ib3
-instance_shared_dev /db2instance
-tbdev /dev/hdisk1
-u db2sdfel
db2sdin1
```

The *test1-ib0* is the cluster interconnect netname for the *test1* member host, the *test2-ib0*, *test2-ib1*, *test2-ib2*, and *test2-ib3* are cluster interconnect netnames for the *test2* host. This example sets up the cluster caching facility on the *test2* host and a DB2 member on the *test1* host for the instance *db2sdin1* associated with the DB2 installation under */DB2DIR*. The cluster interconnect netnames *test1-ib0* and *test2-ib0* must be on the same IP subnet. The *db2sdfel* value is the fenced user and the *db2sdin1* value is the DB2 pureScale instance owner.

What to do next

After creating your DB2 pureScale instance with the **db2icrt** command:

1. Add a second cluster caching facility.
2. Add any additional communication adapter ports to the CFs. This step is not required if all cluster interconnect netnames were specified during instance creation and while adding the second CF.
3. Add any additional members to the instance.
4. Start the instance.

You can add more hosts as members or as a second cluster caching facility by using the **db2isct** command or the **db2iupdt -add** command. Additional communication adapter ports can be added to CFs with either the **db2isct** command or the **db2iupdt -update** command. You can also remove a member or cluster caching facility from your instance by using the **db2iupdt -drop** command.

Adding communication adapter ports to a CF or member

Cluster caching facilities (CFs) and members support multiple communication adapter ports. Using multiple communication adapter ports provide greater scalability and availability. This task describes how to add a communication adapter port.

Before you begin

- Read the Chapter 12, “Network topology configuration considerations for DB2 pureScale environments,” on page 145.

- Log in as a user with root authority on a host in the DB2 pureScale instance.

About this task

Adding communication adapter ports helps achieve the following goals:

- Increase bandwidth to handle more requests by using more than one communication adapter port.
- Enable the adoption of network topologies with multiple switches.

Restrictions

1. A maximum of four communication adapter ports can be used on each CF or member. Communication adapter ports with one or two ports are supported.
2. On supported AIX operating systems, a CF or member with more than one communication adapter port must be on its own physical host (or LPAR on AIX).

Procedure

To enable additional communication adapter ports:

1. Attach an extra interconnect cable from the switch to an open communication adapter port on the CF or member.
2. Configure the IP address and the associated cluster interconnect netname.
3. Repeat steps 1 - 2 for any additional ports you plan to use.
4. To update the CF or member with the additional cluster interconnect netname configured in the preceding step, run the **db2iupdt -update** command on each CF or member with the additional port.

For example, to update the CF:

```
db2iupdt -update -cf <cfHostname> -cfnet <cfnet0,cfnet1> <instanceName>
```

To update a member:

```
db2iupdt -update -m <memberHostname> -mnet <membernet0,membernet1> <instanceName>
```

Results

The cluster caching facility or member now communicates through each communication adapter port by using the registered cluster interconnect netnames.

Example

Updating a CF and member to use an additional communication adapter port on an InfiniBand network.

1. Before updating the CF or member, db2nodes.cfg contains:

```
0   memberhost0   0   membernet0
128 cfhost0       0   cfnet0
```

Note: Do not modify db2nodes.cfg directly.

2. Update the CF and member by running these commands:

```
db2iupdt -update -cf cfhost0 -cfnet cfnet0,cfnet1 <instanceName>
```

```
db2iupdt -update -m memberhost0 -mnet membernet0,membernet1 <instanceName>
```

After the updates, the db2nodes.cfg contains:

```
0   memberhost0   0   membernet0,membernet1
128 cfhost0       0   cfnet0,cfnet1
```

What to do next

Repeat the same procedure on the secondary CF or other members.

Adding new members or an additional cluster caching facility

You can use the **db2iupdt** or the **db2isetup** command to add new DB2 members or an additional cluster caching facility (CF) to your DB2 pureScale instance.

Before you begin

When using the **db2iupdt** or the **db2isetup** command to add a member or a CF, ensure that you review the prerequisites and the preinstallation checklist so that your instance and hosts are compliant.

You can alternatively use the **db2isetup** command to extend the DB2 pureScale instance to multiple hosts (as opposed to the **db2iupdt -add** command which can extend the instance only one member or one CF at a time).

Topology changes

If you change the member topology, for example by adding or dropping a member, you must take an offline backup before you can access the database. If you attempt to access the database before taking an offline backup, the database is placed in a backup pending state.

You can add multiple members or drop multiple members without having to take a backup after each change. For example, if you add three members, you have to take a backup only after you have completed all add operations. However, if you add two members and then drop a member, or if you drop two members and then add a member, you must take a backup before you can perform any additional member topology changes.

Procedure

1. Stop the instance and all process that are running for the instance.
2. You can use either the **db2iupdt** command or the **db2isetup** command:
 - **Using the db2isetup command:**
 - a. Type the **db2isetup** command to launch the DB2 Instance Setup wizard.
 - b. Proceed to the “Setup a DB2 instance” panel. Select the **Configure new function for an existing DB2 instance** option and click **Next**.
 - c. On the Host List panel, add the hosts that will be part of the DB2 pureScale instance. Hosts are validated as they are added to the list. The existing hosts for the instance are automatically loaded into the host list, but you cannot perform any action on them. After you add the hosts you want, click **Next**.
 - d. Click **Finish** to perform the action.
 - **Using the db2iupdt command:**

Run the following command from a host that is already part of the DB2 pureScale instance:

```
db2iupdt
-add [-m <host1> -mnet <netname>] | [-cf <host2> -cfnet <netname>]
-i <db2ssid>
instance
```

- The netname represents the cluster interconnect netname which corresponds to the communication adapter port used for high speed communication between DB2 members and cluster caching facilities in the DB2 instance.
- For DB2 managed GPFS, the `db2sshid` is the non-root user ID required to SSH between hosts in the GPFS domain.
- For the host name, you need to use the same host name format as the return value of the **hostname** command

Use the `-m` option for a member or the `-cf` option for a cluster caching facility to specify which type of host you want to add. You must select one, but not both. If you want to add multiple hosts to the existing instance at one time, use the **db2isetup** command.

Example

For example, the following **db2iupdt** command adds the host named `member1` with the cluster interconnect netname `member1-ib0` as a member to the DB2 pureScale instance:

```
db2iupdt -add -m member1 -mnet member1-ib0 -i db2ssh1 db2sdin1
```

Similarly, the following **db2iupdt** command adds the host named `cf2` as secondary cluster caching facility with four cluster interconnect netnames to the DB2 pureScale instance `db2sdin1`:

```
db2iupdt -add -cf cf2 -cfnet cf2-ib0,cf2-ib1,cf2-ib2,cf2-ib3 -i db2ssh1 db2sdin1
```

The four cluster interconnect netnames are `cf2-ib0`, `cf2-ib1`, `cf2-ib2`, and `cf2-ib3`.

Updating an instance to a DB2 pureScale instance with the **db2iupdt** command

The **db2iupdt** command can be used to update an instance other than a DB2 pureScale instance to a DB2 pureScale instance.

Before you begin

To run some of these steps, you need root authority.

Before running the **db2iupdt** command, you must first stop the instance and all processes that are running for the instance.

Ensure that you review the prerequisites and the preinstallation checklist so that your instance and hosts are compliant.

About this task

If you have the DB2 pureScale Feature product installed, the **db2iupdt** command or the **db2isetup** command can be used to change an existing instance other than a DB2 pureScale instance to a DB2 pureScale instance (dsf-type instance). In this case, the existing DB2 instance must be at the Version 10 level.

Procedure

1. As the instance owner, stop the instance and all process that are running for the instance.
2. As root, you can use the **db2iupdt** command or the **db2isetup** command:

- **Using the `db2isetaup` command:**
 - a. Type the **db2isetaup** command to launch the DB2 Instance Setup wizard.
 - b. Proceed to the "Set up a DB2 instance" panel. Select the **Configure new function for an existing DB2 instance** option. Click **Next**.
 - c. On the Set up a DB2 Cluster File System panel, the shared disk partition device path and the tiebreaker device path are already complete.
 - d. On the Host List panel, add hosts to the DB2 pureScale instance configuration. Hosts are validated as they are added to the list. After you add the hosts you want, click **Next**.
 - e. Click **Finish** to perform the action.

- **Using the `db2iupdt` command:**

Run the **db2iupdt** command as follows:

```
DB2DIR/instance/db2iupdt -d
-cf host2
-cfnet host2-ib0
-m host1
-mnet host1-ib0
-tbdev /dev/hdisk2
-u db2fenc1
db2inst1
```

where *DB2DIR* represents the installation location of your DB2 copy.

This command uses `/dev/hdisk1` path to create a shared file system to store instance shared files (where the database data is stored) and sets up `/dev/hdisk2` file path as the shared device path of the tiebreaker device. The value of the **-tbdev** parameter must be different from the value of the **-instance_shared_dev** parameter.

What to do next

Ensure that your instance is properly set up. See the "Verifying your DB2 pureScale Feature installation and setup" topic for more details.

Applying fix packs

It is recommended that you keep your DB2 database environment running at the latest fix pack level to ensure problem-free operation. To install a fix pack successfully, perform all of the necessary preinstallation and post-installation tasks.

About this task

A DB2 fix pack contains updates and fixes for problems (Authorized Program Analysis Reports, or "APARs") found during testing at IBM, as well as fixes for problems reported by customers. The `APARLIST.TXT` file describes the fixes contained in each fix pack and it is available for download at `ftp://ftp.software.ibm.com/ps/products/db2/fixes/english-us/aparlist/`.

Fix packs are cumulative. This means that the latest fix pack for any given version of DB2 database contains all of the updates from previous fix packs for the same version of DB2 database.

The fix pack images available are:

- A single server image.

The single server image contains the new and updated code required for all DB2 database server products and the IBM Data Server Client. If more than one DB2 database server product is installed in a single location, the DB2 database server fix pack applies maintenance code updates to all the installed DB2 database server products. The Data Server Client fix pack is contained within the one DB2 database server fix pack (namely the fix pack that can service any one of the following database server products: DB2 Enterprise Server Edition, DB2 Workgroup Server Edition, DB2 Express Edition, DB2 Connect Enterprise Edition, DB2 Connect Application Server Edition, DB2 Connect Unlimited Edition for zSeries, and DB2 Connect Unlimited Edition for i5/OS®). You can use the DB2 database server fix pack to upgrade a Data Server Client.

A single server image can also be used to install any of the DB2 database server products, at a particular fix pack level, with a DB2 try and buy license by default.

The single server fix pack image contains DB2 try-and-buy licenses for all DB2 server products. When you select a new DB2 server product to install or a previously installed DB2 server product to update, the try-and-buy licenses are installed. The try-and-buy licenses do not affect any valid licenses already installed in the same DB2 installation path. Regarding DB2 Connect server products, if you run the **db2licm -l** command to query valid licenses, the try-and-buy license for DB2 Connect server product might display as an invalid license. However, if you do not need to use the DB2 Connect functionality, you can ignore the report. To remove the try-and-buy license for DB2 Connect server, use the **db2licm** command.

- A fix pack for each of the other DB2 database products.

Use this fix pack only if you only have non-server database products or add-on products installed. For example, IBM Data Server Runtime Client.

Do not use this type of fix pack if the installed DB2 database products are only DB2 database server products or a Data Server Client. Instead, use the single server image fix pack.

For Windows platforms, if you have more than one DB2 database product (which includes at least one product that is not a Data Server Client or a DB2 database server) installed in a single DB2 copy, you must download and uncompress all of the corresponding product-specific fix packs before starting the fix pack installation process.

- A universal fix pack.

The universal fix pack services installations where more than one DB2 database product has been installed.

The universal fix pack is not needed if the installed DB2 database products are only DB2 database server products or a Data Server Client. In this case, the single server image fix pack should be used.

On Linux or UNIX operating systems, if national languages have been installed, you also require a separate national language fix pack. The national language fix pack can not be installed alone. A universal or product-specific fix pack must be applied at the same time and they must both be at the same fix pack level. For example, if you are applying a universal fix pack to non-English DB2 database products on Linux or UNIX, you must apply both the universal fix pack and the national language fix pack to update the DB2 database products.

Restrictions

- A DB2 Version 10.1 fix pack can only be applied to DB2 Version 10.1 general availability (GA) or DB2 Version 10.1 fix pack copies.

- All DB2 instances, DAS, and applications related to the DB2 copy being updated must be stopped before installing a fix pack.
- In a partitioned database environment, before installing the fix pack, you must stop the database manager on all database partition servers. You must install the fix pack on the instance-owning database partition server and all other database partition servers. All computers participating in the instance must be updated to the same fix pack level.
- On Linux or UNIX operating systems:
 - If you have DB2 database products on a Network File System (NFS), you must ensure the following applications are stopped completely before installing the fix pack: all instances, the DB2 administration server (DAS), interprocess communications (IPC), and applications on other machines using the same NFS mounted installation.
 - If the system commands **fuser** or **lsof** are not available, the **installFixPack** command cannot detect loaded DB2 database files. You must ensure no DB2 files are loaded and provide an override option to install the fix pack. On UNIX, the **fuser** command is required to check for loaded files. On Linux, either the **fuser** command or **lsof** command is required.
For details on the override option, see the **installFixPack** command.
- On client applications, after a fix pack has been applied, to perform autobind of applications, the user must have bind authority.
- Installation of a DB2 fix pack will not service IBM Data Studio.

Procedure

To install a fix pack:

1. Check fix pack prerequisites.
2. Perform the necessary tasks before installing a fix pack.
3. Choose a fix pack installation method and install the fix pack.
4. Perform the necessary tasks after installing the fix pack.
5. Apply the appropriate DB2 database product license.
If a previously licensed copy of a DB2 database server product does not already exist on the machine, a single server fix pack image can be used to install any of the DB2 database server products. In this case, the DB2 database product installed is treated as a try and buy license, and will stop working after a 90 day trial period unless you upgrade the try and buy license.

What to do next

Check the log file for any post-installation steps, or error messages and recommended actions.

For non-root installations on Linux or UNIX, root-based features (such as High Availability and operating system-based authentication) can be enabled using the **db2rfe** command. If root-based features were enabled after installing your DB2 database product, you must rerun the **db2rfe** command each time a fix pack is applied in order to re-enable those features.

If you have multiple DB2 copies on the same system, those copies can be at different version and fix pack levels. If you want to apply a fix pack to one or more DB2 copies, you must install the fix pack on those DB2 copies one by one.

Updating an instance to a higher level within a release using the **db2iupdt** command

The **db2iupdt** command can be used to update an instance to a higher level within a release.

Before you begin

Before running the **db2iupdt** command, you must first stop the instance and all processes that are running for the instance. When using this command, ensure you have reviewed the prerequisites and the pre-installation checklist so that your instance and hosts are compliant.

About this task

The **db2iupdt** command can be issued against instances of the same version that are associated with the same or a different DB2 copy. In all cases, it will update the instance so that it runs against the code located in the same DB2 copy as where you issued the **db2iupdt** command. Use this command to:

- Install a new DB2 database product or feature to the DB2 copy associated to the DB2 instance.
- Update a DB2 instance from one DB2 copy to another DB2 copy of the same version of DB2 database product.

You must run this command on each host in the DB2 pureScale instance.

Procedure

1. Log in as root.
2. Stop the instance and all process that are running for the instance.
3. Run the **db2iupdt** command as follows to update the `db2inst1` instance:

```
/opt/IBM/db2/copy1/instance/db2iupdt -d db2inst1
```

If this command is run from a DB2 pureScale Feature copy, the existing `db2inst1` must have an instance type of `dsf`.

Enable or disable remote root login

When necessary, remote login for the root account can be enabled or disabled as required.

About this task

After installation, upgrade or host maintenance operations, remote login for the root account can be disabled.

Procedure

1. To enable remote root login, enter the following command:

```
/etc/ssh/sshd_config:  
PermitRootLogin yes #enabled
```
2. To disable remote root login, enter the following command:

```
/etc/ssh/sshd_config:  
PermitRootLogin no #disabled
```

Results

For the change to take effect, the ssh daemon must be restarted:

```
/etc/init.d/sshd restart
```

Changing the db2sshid user ID to a different user

For DB2 managed GPFS filesystems, the non-root user ID (db2sshid) is used to establish a secure shell (SSH) network protocol between a local host and a remote host. This user ID defaults to be the instance owner, but can be changed.

Procedure

To change the db2sshid to a different user:

1. Create the new user on every host in the GPFS domain with the same UID and GID.
2. Setup SSH for this new userid. To obtain and setup Open Secure Shell (OpenSSH), see “Installing and setting up OpenSSH” on page 41.
3. On each host, update the configuration file db2ssh.cfg to the new user ID.

Set the new user ID on each host:

```
db2locssh set_db2sshid <newuserID>
```

where <newuserID> specifies the new user ID.

Verify the user ID is changed:

```
db2locssh display_config
```

Results

The db2ssh.cfg configuration file is updated to the new db2sshid user ID.

Re-adding a deleted db2sshid user ID

For DB2 managed GPFS filesystems, the non-root user ID (db2sshid) is used to establish a secure shell (SSH) network protocol between a local host and a remote host. If accidentally deleted on one or more hosts, the user ID must be recreated on all hosts where it was deleted.

Procedure

To re-add a deleted db2sshid user ID on a host:

1. Recreate the user ID with the same UID and GID as they exist on the other hosts in the GPFS domain.
2. Setup SSH for this new userid. To obtain and setup Open Secure Shell (OpenSSH), see “Installing and setting up OpenSSH” on page 41
3. Update the db2sshid user ID setting in the db2ssh.cfg configuration file:

```
db2locssh set_db2sshid <db2sshID>
```

where <db2sshID> is the db2sshid user ID.

4. Verify the db2sshid user ID setting is updated:

```
db2locssh display_config
```

What to do next

The user ID must be recreated on all hosts where it was deleted.

Chapter 19. Creating a shared file system

Once you have created a DB2 pureScale instance, you must create the file systems for the database directories and log files, before you create a database.

Before you begin

You must have a list of storage paths that are connected to the hosts in the cluster. The disk cannot be currently used by any other file system on any host with connectivity to this disk. To perform this task, you must either own the disks and have read and write access to them, or be the DB2 cluster services administrator.

Procedure

1. Use the **db2cluster** command to create the file system for the database:

```
db2cluster -create -filesystem database-filesystem-name -disk storage-paths \
          -mount mount-point
```

Note: The file system name in the command is created under a root file system name of /db2fs unless a mount point is provided.

2. Use the **db2cluster** command to create the file system for the log files:

```
db2cluster -create -filesystem logfile-filesystem-name -disk storage-paths \
          -mount mount-point
```

3. Optional: Check that the file systems were mounted on their hosts:

```
mount | grep filesystem-name
```

This command can be run from any host in the DB2 pureScale instance.

4. Check the amount of space on the file system:

```
df filesystem-directory
```

This command can be run from any host in the DB2 pureScale instance. For a large file system, it might take a few minutes for all of the storage to register with the operating system. Reissue this command after a few minutes if the amount of space is less than expected immediately after the file system creation command completes.

Results

Once the file system has been created by the **db2cluster** command, the file system is mounted and available for usage on the local host. The instance owner has immediate read and write access to the file system. If a root user created the file system, there are no disk ownership checks.

Example

A DBA installed the DB2 pureScale Feature and created the DB2 pureScale instance using the **db2setup** GUI. The four available disks are hdisk3, hdisk4, hdisk5, and hdisk6, as shown by a **ls** command:

```
ls -l /dev/hdisk3 /dev/hdisk4 /dev/hdisk5 /dev/hdisk6
```

```
brw----- 1 db2inst1 system      17, 12 Nov 11 10:39 /dev/hdisk3
```

```
brw----- 1 db2inst1 system      17, 14 Nov 11 10:39 /dev/hdisk4
brw----- 1 db2inst1 system      17, 12 Nov 11 10:39 /dev/hdisk5
brw----- 1 db2inst1 system      17, 14 Nov 11 10:39 /dev/hdisk6
```

The DBA creates the file systems for the database:

```
db2cluster -create -filesystem mydb2database -disk /dev/hdisk3,/dev/hdisk4 \
-mount /db2fs/mydb2database
```

Next, the DBA creates the file systems for the log files:

```
db2cluster -create -filesystem mydb2logs -disk /dev/hdisk5,/dev/hdisk6 \
-mount /db2fs/mydb2logs
```

Now that the db2database and db2logs file systems are created, the DBA checks to see whether the file system for the database is mounted:

```
mount | grep mydb2database
```

```
/dev/mydb2database /db2fs/mydb2database mmfs rw,mtime,atime,dev=mydb2database
```

If the file system was successfully mounted, the DBA can check the amount of space on the file system:

```
df /db2fs/mydb2database
```

```
Filesystem          512-blocks Free      %Used Iused  %Iused Mounted on
/dev/mydb2database 1048576000 739581952   30%   263665    52% /db2fs/mydb2database
```

Finally, the DBA creates the database:

```
db2 create database testdb1 on /db2fs/mydb2database dbpath on /db2fs/mydb2logs
```

What to do next

Once the file systems for the database directory and log directories are created, you can create a database. If the root user created the file system, the root user is required to create a directory (or change the ownership of the current mount point directory), so the instance owner has access to the file system.

Deleting a shared file system

If you no longer require a particular file system, you can delete it from the DB2 cluster services file system list by using the **db2cluster** command.

Before you begin

The file system must be completely empty before it can be deleted.

To perform this task, you must be either the user ID that created the file system, or the DB2 cluster services administrator.

About this task

Once the file system is deleted, any disks that the file system used are released back to the operating system.

Procedure

To delete a shared file system, use the **db2cluster** command:

```
db2cluster -delete -filesystem filesystem-name
```

Retrieving file system information

You can use the **db2cluster** command to obtain information about file systems and the disks that comprise each file system.

The **db2cluster** command provides an easy way of finding out more information about:

The file systems managed by DB2 cluster services

To list the file systems, use this **db2cluster** command:

```
db2cluster -cfs -list -filesystem
```

An example of the output from this command:

FILE SYSTEM NAME	MOUNT_POINT
-----	-----
prodgpfs4	/prodgpfs4

The disks that are used by a file system

To list the disks, issue the **db2cluster** command:

```
db2cluster -list -filesystem filesystem-name -disk
```

For example, for a file system called *prodgpfs4*, you use this command:

```
db2cluster -list -filesystem prodgpfs4 -disk
```

An example of the output from this command:

PATH ON LOCAL HOST	OTHER KNOWN PATHS
-----	-----
/dev/dm-0	
/dev/dm-1	
/dev/dm-2	
/dev/dm-3	
/dev/dm-4	

PATH ON LOCAL HOST might vary from host to host. An asterisk beside a disk indicates that the disk is a tie-breaker disk. It cannot be removed from the file system unless another disk is assigned as the tiebreaker.

OTHER KNOWN PATHS shows if a disk appears with a different device path name on another host.

The configuration of a file system

To list the configuration of a file system, issue this **db2cluster** command:

```
db2cluster -cfs -list -filesystem filesystem-name -configuration
```

For example, for a file system called *prodgpfs4*, you issue this command:

```
db2cluster -cfs -list -filesystem prodgpfs4 -configuration
```

An example of the output from this command:

```
prodgpfs4 options.
OPTION                                VALUE

minFragmentSize                      32768
inodeSize                            512
indirectBlockSize                    32768
defaultMetadataReplicas              1
maxMetadataReplicas                  2
defaultDataReplicas                  1
maxDataReplicas                      2
blockAllocationType                   cluster
fileLockingSemantics                  nfs4
```

ACLSemantics	all
estimatedAverageFileSize	1048576
numNodes	32
blockSize	1048576
quotasEnforced	none
defaultQuotasEnabled	none
maxNumberOfInodes	2048006
filesystemVersion	10.01 (3.2.1.5)
filesystemVersionLocal	10.01 (3.2.1.5)
filesystemVersionManager	10.01 (3.2.1.5)
filesystemVersionOriginal	10.01 (3.2.1.5)
filesystemHighestSupported	10.01 (3.2.1.5)
supportForLargeLUNs	yes
DMAPIEnabled	no
logfileSize	4194304
exactMtime	yes
suppressAtime	no
strictReplication	whenpossible
storagePools	system
disks	nsd1GPFS;nsd2GPFS;nsd3GPFS;nsd4GPFS;nsd5GPFS
automaticMountOption	yes
additionalMountOptions	none
defaultMountPoint	/prodgpfs4

Chapter 20. Configuring a GDPC environment

The topics in this section detail how to configure a geographically dispersed DB2 pureScale cluster (GDPC) environment.

Before you begin

Ensure the following are completed:

- Sites A, B and C can communicate with each other through reliable TCP/IP links.
- All DB2 pureScale installation prerequisites have been satisfied across all hosts to be used in the cluster.
- Sites A and B are connected via a WAN or dark fiber with distance range extenders, with a single high speed interconnect subnet configured across the sites A and B.
- Sites A and B each have a local SAN controller, and the SAN is zoned such that LUNs used for the DB2 pureScale instance are directly accessible from both sites. A one-to-one mapping between LUNs is required across sites so each LUN on site A has a corresponding equally sized LUN on site B.

For this example, the following hardware configurations are used:

- Site A: Hosts hostA1, hostA2, hostA3
- Site B: Hosts hostB1, hostB2, hostB3
- Site C: Host T

Equal sized LUNs have been provisioned on storage at sites A and B, and all LUNs are accessible by all hosts at sites A & B.

LUNs on disks located at Site A are:

```
/dev/hdiskA1  
/dev/hdiskA2  
/dev/hdiskA3  
/dev/hdiskA4  
/dev/hdiskA5  
/dev/hdiskA6  
/dev/hdiskA7
```

Where /dev/hdiskA1 is used for the instance shared file system; /dev/hdiskA2 is used for the DB log file system; and /dev/hdiskA3, /dev/hdiskA4, /dev/hdiskA5, /dev/hdiskA6, and /dev/hdiskA7 are used for the DB data file system.

LUNs on disks located at Site B are:

```
/dev/hdiskB1  
/dev/hdiskB2  
/dev/hdiskB3  
/dev/hdiskB4  
/dev/hdiskB5  
/dev/hdiskB6  
/dev/hdiskB7
```

Where /dev/hdiskB1 is used for the instance shared file system; /dev/hdiskB2 is used for the DB log file system; and /dev/hdiskB3, /dev/hdiskB4, /dev/hdiskB5, /dev/hdiskB6, and /dev/hdiskB7 are used for the DB data file system.

LUNs on disks located at Site C are as follows. These disks can be as small as 50 MB volumes.

```
/dev/hdiskC1  
/dev/hdiskC2  
/dev/hdiskC3
```

About this task

The geographically dispersed DB2 pureScale cluster (GDPC) will be setup in the following way:

- Database MYDB is to be created on instance *db2inst1*.
- *db2inst1* will have three file systems:
 - *logfs* for transaction logs and database metadata for MYDB.
 - *datafs* for database containers for MYDB.
 - *db2fs1* for the shared file system for the instance

Each command is specified with the following format:

```
uid@host> command
```

Where *uid* is the user ID that executes the command, *host* is where the command should be executed, and *command* is the command to execute.

Getting the cluster installed and running in a GDPC environment

There are procedures to be followed in order to get the geographically dispersed DB2 pureScale cluster (GDPC) installed, and up and running.

Procedure

1. Install the DB2 pureScale feature on sites A and B.

Install the DB2 pureScale Feature on sites A and B using the **db2setup** command. Using the Advanced Configuration menu, designate hostA3 and hostB3 as the CFs and (optionally) one of the two to be the preferred primary CF. On site A, designate hostA1, hostA2, hostB1, and hostB2 as members where hostB1 is the shared disk member and hostB2 is the tiebreaker member. During install the tiebreaker disk must be set up using one of the LUNs. This is temporary and can be changed later. For the example below, an option is to use hdiskA2.

The file system that the **db2setup** command creates for the shared instance metadata is initially a non-replicated GPFS file system. This is converted later to a replicated file system across the sites.

2. Updating majority quorum and SCSI-3 PR settings

Following **db2setup**, the tiebreaker setting might need to be updated to use Majority Node Set. Query the current tiebreaker device using the following command:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /home/db2inst1/sqllib/bin/db2cluster -cm -list  
-tiebreaker
```

If the output from the above command does not specify 'Majority Node Set' as the quorum device, then it must now be updated as follows:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /home/db2inst1/sqllib/bin/db2cluster -cm -set -tiebreaker -majority  
Configuring quorum device for domain 'db2domain_20110224005525' ...  
Configuring quorum device for domain 'db2domain_20110224005525' was successful.
```

After updating the tiebreaker device, verify the setting using the following command and compare it to the expected output:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /home/db2inst1/sqllib/bin/db2cluster -cm -list
-tiebreaker
The current quorum device is of type Majority Node Set.
```

Note: If the third site does not have direct access to the disks on site A or B then SCSI-3 PR must be disabled.

Check to see if SCSI-3 PR is enabled. In the sample output, pr=yes indicates SCSI-3 PR is enabled:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmfsnsd -X
```

```
Disk name NSD volume ID Device Devtype Node name Remarks
-----
gpfs1nsd 091A33584D65F2F6 /dev/hdiskA1 hdisk hostA1 pr=yes
```

If your disks do not support SCSI-3 PR or you choose to disable it, run these commands:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> su - db2inst1
db2inst1@hostA1:/home/db2inst1> db2stop force
02/24/2011 01:24:16 0 0 SQL1064N DB2STOP processing was successful.
02/24/2011 01:24:19 1 0 SQL1064N DB2STOP processing was successful.
02/24/2011 01:24:21 3 0 SQL1064N DB2STOP processing was successful.
02/24/2011 01:24:22 2 0 SQL1064N DB2STOP processing was successful.
SQL1064N DB2STOP processing was successful.
db2inst1@hostA1:/home/db2inst1> exit
root@hostA1:/opt/IBM/db2/V10.1/bin> /home/db2inst1/sqllib/bin/db2cluster -cfs -stop -all
All specified hosts have been stopped successfully.
```

Verify that GPFS is stopped on all hosts:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmgetstate -a
Node number Node name GPFS state
-----
```

```
1 hostA1 down
2 hostA2 down
3 hostA3 down
4 hostB1 down
5 hostB2 down
6 hostB3 down
```

GPFS is down, so disable SCSI-3 PR:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmchconfig usePersistentReserve=no
Verifying GPFS is stopped on all nodes ...
mmchconfig: Processing the disks on node hostA1.torolab.ibm.com
mmchconfig: Processing the disks on node hostA2.torolab.ibm.com
mmchconfig: Processing the disks on node hostA3.torolab.ibm.com
mmchconfig: Processing the disks on node hostB1.torolab.ibm.com
mmchconfig: Processing the disks on node hostB2.torolab.ibm.com
mmchconfig: Processing the disks on node hostB3.torolab.ibm.com
mmchconfig: Command successfully completed
mmchconfig: Propagating the cluster configuration data to all affected nodes. This
is an asynchronous process.
```

Verify that SCSI-3 PR has been disabled (pr=yes is not displayed):

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmfsnsd -X
Disk name NSD volume ID Device Devtype Node name Remarks
-----
gpfs1nsd 091A33584D65F2F6 /dev/hdiskA1 hdisk hostA1
```

Verify that **usePersistentReserve** has been set to no:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmfsconfig
Configuration data for cluster db2cluster_20110224005554.torolab.ibm.com:
-----
clusterName db2cluster_20110224005554.torolab.ibm.com
clusterId 655893150084494058
autoload yes
minReleaseLevel 3.4.0.7
dmapiFileHandleSize 32
maxFilesToCache 10000
pagepool 256M
verifyGpfsReady yes
assertOnStructureError yes
worker1Threads 150
sharedMemLimit 2047M
```

```

usePersistentReserve no
failureDetectionTime 35
leaseRecoveryWait 35
tiebreakerDisks gpfs1nsd
[hostA1]
psspVsd no
adminMode allToAll
File systems in cluster db2cluster_20110224005554.torolab.ibm.com:
-----
/dev/db2fs1

```

3. Increase HostFailureDetectionTime.

HostFailureDetectionTime is increased to a higher value than what would be set on a non-GDPC DB2 pureScale cluster. Changing this value allows for the increased communication lag between sites that is not present in a single-site DB2 pureScale cluster. If unexpected host down events are still triggered due to large inter-site distances, higher parameter values can be used, however this will increase the time required for DB2 pureScale to detect hardware failures or machine reboots, increasing the overall failure recovery time.

```

root@hostA1:/opt/IBM/db2/V10.1/bin> /home/db2inst1/sqllib/bin/db2cluster -cm -list
-hostfailedetectiontime
The host failure detection time is 4 seconds.

```

Change it to 16 seconds and verify.

```

root@hostA1:/opt/IBM/db2/V10.1/bin> /home/db2inst1/sqllib/bin/db2cluster -cm -set -option hostfailedetectiontime
-value 16
The host failure detection time has been set to 16 seconds.

root@hostA1:/opt/IBM/db2/V10.1/bin> /home/db2inst1/sqllib/bin/db2cluster -cm -list
-hostfailedetectiontime
The host failure detection time is 16 seconds.

```

4. Add tiebreaker host into cluster.

The tiebreaker host provides cluster quorum, ensuring that during normal operation, the cluster contains an odd number of hosts. In case of a network outage between sites, only the site which can communicate with tiebreaker host T will gain cluster quorum.

Follow the steps in the Installation prerequisites section of the DB2 documentation to install the appropriate uDAPL level on the tiebreaker host. Then install DB2 software on the tiebreaker host:

```

root@T:/path containing db2_install. /db2_install

Default directory for installation of products - /opt/IBM/db2/V9.8
*****
Do you want to choose a different directory to install [yes/no] ?
no

Specify one of the following keywords to install DB2 products.
ESE_DSF

Enter "help" to redisplay product names.
Enter "quit" to exit.
*****
ESE_DSF

DB2 installation is being initialized.

Total number of tasks to be performed: 46
Total estimated time for all tasks to be performed: 2850 second(s)

Task #1 start
...

Task #46 end

The execution completed successfully.
For more information see the DB2 installation log at /tmp/db2_install.log.nnnnnnnn.

```

Change the GPFS quorum type for the cluster to majority node set and verify:

```

root@hostA1:/opt/IBM/db2/V10.1/bin> /home/db2inst1/sqllib/bin/db2cluster -cfs -set -tiebreaker -majority
root@hostA1:/opt/IBM/db2/V10.1/bin> /home/db2inst1/sqllib/bin/db2cluster -cfs -list
-tiebreaker
The current quorum device is of type Majority Node Set.

```

Ensure SSH is setup for the user, in our example *db2inst1*, on the tiebreaker host T. This user should be the same *db2sshid* user set during the installation on site A and site B. To check what user was used, run the following command on hostA:

```
root@hostA1>/var/db2/db2ssh/db2locssh display_config
```

```
version = 1
time_delta = 20 second(s)
debug_level = 2
db2sshid = db2inst1
gdkit_path = /opt/IBM/db2/V10.1/lib64/gskit/
fips_mode = on
```

Run the following commands as root on host T:

- Create the configuration file:
`/var/db2/db2ssh/db2locssh reset_config`
- Set the GSKit path:
`/var/db2/db2ssh/db2locssh set_gskit_path /opt/IBM/db2/V10.1/lib64/gskit/`
- Set the *db2ssh* to *db2sshid* (*db2sshid* might be a new ID or the instance owner):
`/var/db2/db2ssh/db2locssh set_db2sshid db2inst1`
- Verify the setting:
`root@T>/var/db2/db2ssh/db2locssh display_config`

```
version = 1
time_delta = 20 second(s)
debug_level = 2
db2sshid = db2inst1
gdkit_path = /opt/IBM/db2/V10.1/lib64/gskit/
fips_mode = on
```

- Generate a private/public key pair:
`/var/db2/db2ssh/db2locssh generate_keys`

Perform the key exchange for every host in the cluster and perform a test:

```
root@T>/var/db2/db2ssh/db2locssh root@hostA1 hostname
hostA1
```

```
root@hostA1>/var/db2/db2ssh/db2locssh root@T hostname
T
```

Add the tiebreaker host to the RSCT cluster:

```
root@T> preprnode hostA1 hostA2 hostB1 hostB2 hostA3 hostB3

root@hostA1:/opt/IBM/db2/V10.1/bin> lsrrnode
Name OpState RSCTVersion
hostB2 Online 3.1.2.2
hostB3 Online 3.1.2.2
hostA3 Online 3.1.2.2
hostB1 Online 3.1.2.2
hostA2 Online 3.1.2.2
hostA1 Online 3.1.2.2
root@hostA1:/opt/IBM/db2/V10.1/bin> /home/db2inst1/sqllib/bin/db2cluster -cm -add -host T
Adding node 'T' to the cluster ...
Trace spooling could not be enabled on the local host.
Adding node 'T' to the cluster was successful.
```

Verify that the tiebreaker host has been added to the RSCT cluster:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> lsrrnode
Name OpState RSCTVersion
T Online 3.1.2.2
hostB3 Online 3.1.2.2
hostB2 Online 3.1.2.2
hostB1 Online 3.1.2.2
hostA3 Online 3.1.2.2
hostA2 Online 3.1.2.2
hostA1 Online 3.1.2.2
```

Add the tiebreaker host to the GPFS cluster. To mark this host as a quorum client ensure it never runs as a file system manager, token manager, or other role, this is done directly with the GPFS **mmaddnode** command:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmlsnode
GPFS nodeset Node list
-----
db2cluster_20110224005554 hostA1 hostA2 hostA3 hostB1 hostB2 hostB3
root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmaddnode T:quorum-client
```

```
Thu Feb 24 01:49:38 EST 2011: mmaddnode: Processing node T.torolab.ibm.com
mmaddnode: Command successfully completed
mmaddnode: Warning: Not all nodes have proper GPFS license designations.
mmaddnode: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

Verify that the tiebreaker host has been added to the GPFS cluster:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmlsnode
```

```
=====
Warning: |
This cluster contains nodes that do not have a proper GPFS license |
designation. This violates the terms of the GPFS licensing agreement. |
Use the mmchlicense command and assign the appropriate GPFS licenses |
to each of the nodes in the cluster. For more information about GPFS |
license designation, see the Concepts, Planning, and Installation Guide. |
=====
```

```
GPFS nodeset Node list
-----
```

```
db2cluster_20110224005554 hostA1 hostA2 hostA3 hostB1 hostB2 hostB3 T
```

On the tiebreaker host add the GPFS license:

```
root@T:/opt/IBM/db2/V10.1/bin> ./db2cluster -cfs -add -license
```

The license for the shared file system cluster has been successfully added.

Verify the license warning message is gone:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmlsnode
```

```
GPFS nodeset Node list
-----
```

```
db2cluster_20110224005554 hostA1 hostA2 hostA3 hostB1 hostB2 hostB3 T
```

The `/usr/lpp/mmfs/bin/mmlscluster` command can be used to ensure that the tiebreaker host has been added to the GPFS cluster with a designation of “quorum” rather than “quorum-manager”. All other hosts in the cluster should be designated as quorum-manager. If the tiebreaker host is a quorum-manager its status can be changed to client with the `/usr/lpp/mmfs/bin/mmchnode --client -N hostT` command.

If you are not running SCSI-3 PR, and therefore do not require direct access to all disks from the tiebreaker site, you can ensure you will not get false errors in this configuration by following these steps:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmchconfig unmountOnDiskFail=yes -N T
mmchconfig: Command successfully completed
mmchconfig: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmlsconfig
Configuration data for cluster db2cluster_20110224005554.torolab.ibm.com:
```

```
-----
clusterName db2cluster_20110224005554.torolab.ibm.com
clusterId 655893150084494058
autoload yes
minReleaseLevel 3.4.0.7
dmapiFileHandleSize 32
maxFilesToCache 10000
pagepool 256M
verifyGpfsReady yes
assertOnStructureError yes
```



```

worker1Threads 150
sharedMemLimit 2047M
usePersistentReserve no
failureDetectionTime 35
leaseRecoveryWait 35
[T]
unmountOnDiskFail yes
[common]
[hostA1]
psspVsd no
adminMode allToAll
File systems in cluster db2cluster_20110224005554.torolab.ibm.com:
-----
/dev/db2fs1

```

```

root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmgetstate -a

```

```

Node number Node name GPFS state
-----
1 hostA1 down
2 hostA2 down
3 hostA3 down
4 hostB1 down
5 hostB2 down
6 hostB3 down
7 T down

```

```

root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmchconfig unmountOnDiskFail=yes -N T
mmchconfig: Command successfully completed
mmchconfig: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.

```

To verify the line was changed, enter:

```

root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmlsconfig
Configuration data for cluster db2cluster_20110224005554.torolab.ibm.com:

```

```

-----
clusterName db2cluster_20110224005554.torolab.ibm.com
clusterId 655893150084494058
autoload yes
minReleaseLevel 3.4.0.7
dmapiFileHandleSize 32
maxFilesToCache 10000
pagepool 256M
verifyGpfsReady yes
assertOnStructureError yes
worker1Threads 150
sharedMemLimit 2047M
usePersistentReserve no
failureDetectionTime 35
leaseRecoveryWait 35
[T]
unmountOnDiskFail yes
[common]
[hostA1]
psspVsd no
adminMode allToAll
File systems in cluster db2cluster_20110224005554.torolab.ibm.com:
-----
/dev/db2fs1

```

```

root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmgetstate -a

```

```

Node number Node name GPFS state
-----
1 hostA1 down
2 hostA2 down
3 hostA3 down
4 hostB1 down
5 hostB2 down
6 hostB3 down
7 T down

```

To enable GPFS to respond to failures faster, update the **failureDetectionTime** and **leaseRecoveryWait** parameters:

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmchconfig failureDetectionTime=30
Verifying GPFS is stopped on all nodes ...
mmchconfig: Command successfully completed
mmchconfig: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmchconfig leaseRecoveryWait=25
mmchconfig: Command successfully completed
mmchconfig: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

Update the GPFS cluster to use the IB private network to communicate between sites A and B. This enables the clustering software to detect network issues between the sites, and trigger failover accordingly. First, check the subnet for the IB network:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> ping hostA1-ib0
PING hostA1-ib0.torolab.ibm.com (10.1.1.1): 56 data bytes
64 bytes from 10.1.1.1: icmp_seq=0 ttl=255 time=0 ms
```

Note that if the tiebreaker host does not have IB access to the two main sites, update the */etc/hosts* file on the tiebreaker host to map the IB names from the two main sites to the regular Ethernet names for those hosts. This is required by GPFS when creating file systems, since the tiebreaker host does not belong to the IB subnet:

```
root@T:/> cat /etc/hosts
9.26.82.1 hostA1.torolab.ibm.com hostA1
9.26.82.1 hostA1-ib0.torolab.ibm.com hostA1-ib0
9.26.82.2 hostA2.torolab.ibm.com hostA2
9.26.82.2 hostA2-ib0.torolab.ibm.com hostA2-ib0
9.26.82.3 hostA3.torolab.ibm.com hostA3
9.26.82.3 hostA3-ib0.torolab.ibm.com hostA3-ib0
9.26.82.4 hostB1.torolab.ibm.com hostB1
9.26.82.4 hostB1-ib0.torolab.ibm.com hostB1-ib0
9.26.82.5 hostB2.torolab.ibm.com hostB2
9.26.82.5 hostB2-ib0.torolab.ibm.com hostB2-ib0
9.26.82.6 hostB3.torolab.ibm.com hostB3
9.26.82.6 hostB3-ib0.torolab.ibm.com hostB3-ib0
```

In this example, subnet 10.1.1.0 includes all the IP addresses from 10.1.1.0 through 10.1.1.255:

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmchconfig subnets=10.1.1.0
mmchconfig: Command successfully completed
mmchconfig: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

```
root@hostA1:/opt/IBM/db2/V10.1/bin> /usr/lpp/mmfs/bin/mmlsconfig
Configuration data for cluster db2cluster_20110224005554.torolab.ibm.com:
-----
clusterName db2cluster_20110224005554.torolab.ibm.com
clusterId 655893150084494058
autoload yes
minReleaseLevel 3.4.0.7
dmapiFileHandleSize 32
maxFilesToCache 10000
pagepool 256M
verifyGpfsReady yes
assertOnStructureError yes
worker1Threads 150
sharedMemLimit 2047M
usePersistentReserve no
failureDetectionTime 30
leaseRecoveryWait 25
[T]
unmountOnDiskFail yes
[common]
subnets 10.1.1.0
[hostA1]
psspVsd no
adminMode allToAll
```

```
File systems in cluster db2cluster_20110224005554.torolab.ibm.com:
-----
/dev/db2fs1
```

Ensure that each site contains a shared file system configuration server so that the GPFS configuration files will be preserved in case of a disaster on one site. Change the configuration servers so that hostA1 is the primary configuration server and hostB1 is the secondary configuration server:

```
root@hostA1> /usr/lpp/mmfs/bin/mmchcluster -p hostA1 -s hostB1
```

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmlscluster
```

```
GPFS cluster information
```

```
=====
```

```
GPFS cluster name: db2cluster_20110224005554.torolab.ibm.com
```

```
GPFS cluster ID: 655893150084494058
```

```
GPFS UID domain: db2cluster_20110224005554.torolab.ibm.com
```

```
Remote shell command: /usr/bin/ssh
```

```
Remote file copy command: /usr/bin/scp
```

```
GPFS cluster configuration servers:
```

```
-----
```

```
Primary server: hostA1.torolab.ibm.com
```

```
Secondary server: hostB1.torolab.ibm.com
```

Update the RSCT communication groups to disable Loose Source Routing (LSR). When LSR is disabled, RSCT will use daemon routing, which is a more reliable communication method in the event of isolated network failures. First list all the communication groups used by RSCT, and then update each separately:

```
root@hostA1:/> lscomg
Name Sensitivity Period Priority Broadcast SourceRouting NIMPathName NIMParameters Grace MediaType UseForNodeMembership
CG1 4 1.6 1 Yes Yes 60 1 (IP) 1
CG2 4 1.6 1 Yes Yes 60 1 (IP) 1
root@hostA1:/> chcomg -x r CG1
root@hostA1:/> chcomg -x r CG2
root@hostA1:/> lscomg
Name Sensitivity Period Priority Broadcast SourceRouting NIMPathName NIMParameters Grace MediaType UseForNodeMembership
CG1 4 1.6 1 Yes No 60 1 (IP) 1
CG2 4 1.6 1 Yes No 60 1 (IP) 1
```

Note that if at anytime the `db2cluster -cm -delete -domain/create` domain commands are run to recreate the TSA domain, then LSR needs to be disabled again.

For better resilience during Ethernet failures, update `/etc/hosts` on all hosts in the cluster to contain a mapping from each host name to it's IP address (note that from earlier in this step, host T's `/etc/hosts` file will differ from the below, as its `-ib0` hostnames will map to the standard Ethernet hostname). This setting prevents some DB2 Cluster Services monitor commands from hanging in the event that one of the DNS servers at a site has failed:

```
root:/> cat /etc/hosts
10.1.1.1 hostA1-ib0.torolab.ibm.com hostA1-ib0
10.1.1.2 hostA2-ib0.torolab.ibm.com hostA2-ib0
10.1.1.3 hostA3-ib0.torolab.ibm.com hostA3-ib0
10.1.1.4 hostB1-ib0.torolab.ibm.com hostB1-ib0
10.1.1.5 hostB2-ib0.torolab.ibm.com hostB2-ib0
10.1.1.6 hostB3-ib0.torolab.ibm.com hostB3-ib0
9.26.82.1 hostA1.torolab.ibm.com hostA1
9.26.82.2 hostA2.torolab.ibm.com hostA2
9.26.82.3 hostA3.torolab.ibm.com hostA3
9.26.82.4 hostB1.torolab.ibm.com hostB1
9.26.82.5 hostB2.torolab.ibm.com hostB2
9.26.82.6 hostB3.torolab.ibm.com hostB3
9.23.1.12 T
```

Setting up GPFS replication in a GDPC environment

When configuring a geographically dispersed DB2 pureScale cluster (GDPC) environment, you must set up IBM General Parallel File System (GPFS) replication.

Procedure

1. Prepare the **sql1ib_shared** file system for replication.

To enable replication, change the failure group of the non-replicated GPFS file system to 1. This would typically be the failure group on site A.

To permit that operation, restart GPFS and then ensure that the DB2 instance is stopped for each host, so the file system can be unmounted:

```
root@hostA1:/> /home/db2inst1/sql1ib/bin/db2cluster -cfs -start -all
All specified hosts have been started successfully.
```

```
db2inst1@hostA1:/home/db2inst1> db2stop instance on hostA1
SQL1064N DB2STOP processing was successful.
db2inst1@hostA2:/home/db2inst1> db2stop instance on hostA2
SQL1064N DB2STOP processing was successful.
db2inst1@hostA3:/home/db2inst1> db2stop instance on hostA3
SQL1064N DB2STOP processing was successful.
db2inst1@hostB1:/home/db2inst1> db2stop instance on hostB1
SQL1064N DB2STOP processing was successful.
db2inst1@hostB2:/home/db2inst1> db2stop instance on hostB2
SQL1064N DB2STOP processing was successful.
db2inst1@hostB3:/home/db2inst1> db2stop instance on hostB3
SQL1064N DB2STOP processing was successful.
```

To ensure the **sql1ib_shared** file system is cleanly unmounted, the cluster is put in maintenance mode:

```
root@hostA1:/> /home/db2inst1/sql1ib/bin/db2cluster -cm -enter -maintenance -all
Domain 'db2domain_20110224005525' has entered maintenance mode.
```

Changing the failure group of the disk requires us to find out the Network Shared Disk (NSD) name that GPFS assigned to the disk. In the following sample output, the column 'Device' contains the actual device path and the column 'Disk name' contains the NSD name that GPFS assigned to that device.

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmlnsd -m
```

```
Disk name NSD volume ID Device Node name Remarks
gpfs1nsd 091A33584D65F2F6 /dev/hdiskA1 hostA1.torolab.ibm.com
```

Create a file */tmp/nsdAddFGroup.txt* containing a line describing the disk, and which indicates it is part of failure group 1. This file should list all the NSD disks that belong to site A and that will belong to the db2fs1 file system. These disks will then be assigned to the first failure group. In this example, there is just one disk:

```
root@hostA1:/> cat /tmp/nsdAddFGroup.txt
gpfs1nsd:::dataAndMetadata:1
```

```
root@hostA1:/> /home/db2inst1/sql1ib/bin/db2cluster -cfs -list -filesystem
File system NAME MOUNT_POINT
-----
db2fs1 /db2sd_20110224005651
```

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmlsdisk db2fs1 -L
disk driver sector failure holds holds storage
name type size group metadata data status availability disk ID pool remarks
-----
gpfs1nsd nsd 512 -1 yes yes ready up 1 system desc
Number of quorum disks: 1
Read quorum value: 1
Write quorum value: 1
```

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmchdisk db2fs1 change -F /tmp/nsdAddFGroup.txt
Verifying file system configuration information ...
```

mmchdisk: Propagating the cluster configuration data to all affected nodes. This is an asynchronous process.

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmldisk db2fs1 -L
disk driver sector failure holds holds storage
name type size group metadata data status availability disk ID pool remarks
-----
gpfs1nsd nsd 512 1 yes yes ready up 1 system desc
Number of quorum disks: 1
Read quorum value: 1
Write quorum value: 1
Attention: Due to an earlier configuration change the file system
is no longer properly replicated.
```

Note that the disk gpfs1nsd is now assigned to failure group 1 (previously, it was -1)

Change the replication settings for the file system to enable replication:

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmchfs db2fs1 -m 2 -r 2
The desired replication factor exceeds the number of available metadata failure groups.
Allowed, but files will be unreplicated and hence at risk.
Attention: The desired replication factor exceeds the number of available data failure groups in
storage pool system.
This is allowed, but files in this storage pool will not be replicated and will therefore be at risk.
```

Verify that the file system settings have been changed to enable replication:

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmllsfs db2fs1
flag value description
```

```
-----
-f 32768 Minimum fragment size in bytes
-i 512 Inode size in bytes
-I 32768 Indirect block size in bytes
-m 2 Default number of metadata replicas
-M 2 Maximum number of metadata replicas
-r 2 Default number of data replicas
-R 2 Maximum number of data replicas
```

2. Create an affinity between the network shared disk (NSD) and the hosts.

Although some physical storage is local to each site, GPFS does not know which LUN is locally accessible (over the SAN) at each site. However, to indicate to GPFS that it should prefer going to local LUNs for read operations, providing better performance. Create a file */tmp/affinitizensd.txt* to contain a line that indicates the disk is part of site A, and then use **mmchnsd** to create the affinity between the NSD and a site.

```
root@hostA1:/> cat /tmp/affinitizensd.txt
gpfs1nsd:hostA1,hostA2,hostA3
```

Note that the previous step stopped the DB2 pureScale instance and placed the cluster into CM maintenance mode (as oppose to CFS maintenance mode) – this is necessary for the following steps as well. Verify the file system is not mounted. If it is mounted, then unmount it. Unmounting a filesystem can be done with the **db2cluster -cfs -unmount -filesystem *filesystem*** command.

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmllsmount db2fs1
File system db2fs1 is not mounted.
```

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmchnsd -F /tmp/affinitizensd.txt
mmchnsd: Processing disk gpfs1nsd
mmchnsd: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

Verify that the site A computers (hostA*) have become the server hosts for the disk:

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmllsnsd -X
```

```
Disk name NSD volume ID Device Devtype Node name Remarks
gpfs1nsd 091A33584D65F2F6 /dev/hdiskA1 hdisk hostA1.torolab.ibm.com server node
gpfs1nsd 091A33584D65F2F6 /dev/hdiskA1 hdisk hostA2.torolab.ibm.com server node
gpfs1nsd 091A33584D65F2F6 /dev/hdiskA1 hdisk hostA3.torolab.ibm.com server node
```

Restart the cluster:

```
root@hostA1:/> /home/db2inst1/sqllib/bin/db2cluster -cm -exit -maintenance
```

Host 'hostA1' has exited maintenance mode. Domain 'db2domain_20110224005525' has been started.

Verify that the file system has been remounted, and then restart the instance on each computer:

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmlsmount db2fs1
File system db2fs1 is mounted on 6 nodes.
```

```
db2inst1@hostA1:/home/db2inst1> db2start instance on hostA1
SQL1063N DB2START processing was successful.
db2inst1@hostA2:/home/db2inst1> db2start instance on hostA2
SQL1063N DB2START processing was successful.
db2inst1@hostA3:/home/db2inst1> db2start instance on hostA3
SQL1063N DB2START processing was successful.
db2inst1@hostB1:/home/db2inst1> db2start instance on hostB1
SQL1063N DB2START processing was successful.
db2inst1@hostB2:/home/db2inst1> db2start instance on hostB2
SQL1063N DB2START processing was successful.
db2inst1@hostB3:/home/db2inst1> db2start instance on hostB3
SQL1063N DB2START processing was successful.
```

Verify with **db2instance -list** that the host resources are now online for all 6 computers:

```
$ db2instance -list
ID  TYPE      STATE  HOME_HOST  CURRENT_HOST  ALERT  PARTITION_NUMBER  LOGICAL_PORT  NETNAME
---
0   MEMBER    STOPPED hostA1      hostA1        NO      0                  0             hostA1-ib0
1   MEMBER    STOPPED hostA2      hostA2        NO      0                  0             hostA2-ib0
2   MEMBER    STOPPED hostB1      hostB1        NO      0                  0             hostB1-ib0
3   MEMBER    STOPPED hostB2      hostB2        NO      0                  0             hostB2-ib0
128 CF       STOPPED hostA3      hostA3        NO      -                  0             hostA3-ib0
129 CF       STOPPED hostB3      hostB3        NO      -                  0             hostB3-ib0
-----
HOSTNAME  STATE  INSTANCE_STOPPED  ALERT
-----
hostA1    ACTIVE  NO                NO
hostA2    ACTIVE  NO                NO
hostA3    ACTIVE  NO                NO
hostB1    ACTIVE  NO                NO
hostB2    ACTIVE  NO                NO
hostB3    ACTIVE  NO                NO
```

3. Add the replica disk from site B and the file system quorum disk from the tiebreaker site.

Add the replica disk and file system quorum disk to the existing **sqllib_shared** file system. Note that information about the affinity of the LUNs is added to their local hosts.

Create a file **/tmp/nsdfailuregroup2.txt** that describes the replica disk(s) at site B and **/tmp/nsdfailuregroup3.txt** that describes the tiebreaker disk on host T. In the following example **hdiskB1** on Site B will hold the data replica for the **sqllib_shared** file system, while the **hdiskC1** on host T will act as a quorum disk.

```
root@hostA1:/> cat /tmp/nsdfailuregroup2.txt
/dev/hdiskB1:hostB1,hostB2,hostB3::dataAndMetadata:2
root@hostA1:/> /usr/lpp/mmfs/bin/mmcrnsd -F /tmp/nsdfailuregroup2.txt
mmcrnsd: Processing disk hdiskB1
mmcrnsd: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

```
root@T:/> cat /tmp/nsdfailuregroup3.txt
/dev/hdiskC1:T::descOnly:3
```

```
root@T:/> /usr/lpp/mmfs/bin/mmcrnsd -F /tmp/nsdfailuregroup3.txt
mmcrnsd: Processing disk hdiskC1
mmcrnsd: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

Verify that the NSDs have been created with the **mmlnsd** command:

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmlnsd -X
Disk name  NSD volume ID  Device  Devtype  Node name
-----
gpfs100lnsd 091A336D4D674B1E /dev/hdiskB1 hdisk hostA1.torolab.ibm.com
gpfs100lnsd 091A336D4D674B1E /dev/hdiskB1 hdisk hostA2.torolab.ibm.com
gpfs100lnsd 091A336D4D674B1E /dev/hdiskB1 hdisk hostA3.torolab.ibm.com
```

```
gpfs1001nsd 091A33604D674B1E /dev/hdiskB1 hdisk hostB1.torolab.ibm.com server node
gpfs1001nsd 091A33604D674B1E /dev/hdiskB1 hdisk hostB2.torolab.ibm.com server node
gpfs1001nsd 091A33604D674B1E /dev/hdiskB1 hdisk hostB3.torolab.ibm.com server node
gpfs1002nsd 091A33434D674B57 /dev/hdiskC1 hdisk T.torolab.ibm.com server node
gpfs1nsd 091A33584D65F2F6 /dev/hdiskA1 hdisk hostA1.torolab.ibm.com server node
gpfs1nsd 091A33584D65F2F6 /dev/hdiskA1 hdisk hostA2.torolab.ibm.com server node
gpfs1nsd 091A33584D65F2F6 /dev/hdiskA1 hdisk hostA3.torolab.ibm.com server node
gpfs1nsd 091A33584D65F2F6 /dev/hdiskA1 hdisk hostB1.torolab.ibm.com
gpfs1nsd 091A33584D65F2F6 /dev/hdiskA1 hdisk hostB2.torolab.ibm.com
gpfs1nsd 091A33584D65F2F6 /dev/hdiskA1 hdisk hostB3.torolab.ibm.com
```

Add the disk at site B to a file system:

```
root@hostA1:/> /home/db2inst1/sqllib/bin/db2cluster -cfs -add -filesystem db2fs1 -disk /dev/hdiskB1
```

If an **hdisk** with the same name is in use on any host in the cluster you might see the following error:

```
Disk '/dev/hdiskB1' is already in use.
There is a problem with the disks specified in the operation. Check the diagnostic log (db2diag.log or
/tmp/ibm.db2.cluster.*) for more information. Correct the problem and re-issue the command.
A diagnostic log has been saved to '/tmp/ibm.db2.cluster.CJeoEa'.
```

This can happen if, for example, the same disk names are used at the two different sites. If that is the case (and if you are certain that disk is not actually already being used elsewhere), update the `/tmp/nsdfailuregroup2.txt` file similar to the example below, and use the GPFS **mmaddisk** command directly:

```
root@hostA1:/> cat /tmp/nsdfailuregroup2.txt
# /dev/hdiskB1:hostB1,hostB2,hostB3::dataAndMetadata:2
gpfs1001nsd:::dataAndMetadata:2::
```

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmaddisk db2fs1 -F /tmp/nsdfailuregroup2.txt
The following disks of db2fs1 will be formatted on node hostA1:
gpfs1001nsd: size 34603008 KB
Extending Allocation Map
Checking Allocation Map for storage pool 'system'
Completed adding disks to file system db2fs1.
mmaddisk: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

Verify that the disk has been added to the file system with the correct failure group:

```
root@hostA1:/> /usr/lpp/mmfs/bin/mmlsdisk db2fs1 -L
```

```
disk driver sector failure holds holds storage
name type size group metadata data status availability disk ID pool remarks
-----
gpfs1nsd nsd 512 1 yes yes ready up 1 system desc
gpfs1001nsd nsd 512 2 yes yes ready up 2 system desc
Number of quorum disks: 2
Read quorum value: 2
Write quorum value: 2
Attention: Due to an earlier configuration change the file system
is no longer properly replicated.
```

Similarly, add the disk at the tiebreaker site to the file system:

```
root@T:/> cat /tmp/nsdfailuregroup3.txt
# /dev/hdiskC1:T::descOnly:3
gpfs1002nsd:::descOnly:3::
```

```
root@T:/> /usr/lpp/mmfs/bin/mmaddisk db2fs1 -F /tmp/nsdfailuregroup3.txt
```

```
The following disks of db2fs1 will be formatted on node T:
gpfs1002nsd: size 1048576 KB
Extending Allocation Map
Checking Allocation Map for storage pool 'system'
Completed adding disks to file system db2fs1.
mmaddisk: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

Verify that the disk has been added to the file system and to the correct failure group:


```

root@T: /> /usr/lpp/mmfs/bin/mmlsdisk db2fs1 -L

disk driver sector failure holds holds storage
name type size group metadata data status availability disk ID pool remarks
-----
gpfs1nsd nsd 512 1 yes yes ready up 1 system desc
gpfs1001nsd nsd 512 2 yes yes ready up 2 system desc
gpfs1002nsd nsd 512 3 no no ready up 3 system desc
Number of quorum disks: 3
Read quorum value: 2
Write quorum value: 2
Attention: Due to an earlier configuration change the file system
is no longer properly replicated.

```

4. Rebalance the file system to replicate the data on the newly added disks.

```

root@hostA1: /> /usr/lpp/mmfs/bin/mmrestripefs db2fs1 -R

Verify that the message about the file system not being replicated is gone:

root@hostA1: /> /usr/lpp/mmfs/bin/mmlsdisk db2fs1 -L
disk driver sector failure holds holds storage
name type size group metadata data status availability disk ID pool remarks
-----
gpfs1nsd nsd 512 1 yes yes ready up 1 system desc
gpfs1001nsd nsd 512 2 yes yes ready up 2 system desc
gpfs1002nsd nsd 512 3 no no ready up 3 system desc
Number of quorum disks: 3
Read quorum value: 2
Write quorum value: 2

```

At the end of this step, the following is set up:

- A GPFS and RSCT cluster across sites A, B and C
- A tie-breaker host T that is part of the RSCT domain and GPFS cluster but is not part of the DB2 instance.
- A DB2 pureScale cluster spanning sites A and B, with the instance shared metadata **sql1ib_shared** file system being a replicated GPFS file system across sites A and B.

In the example above, the data in **sql1ib_shared** is stored on both /dev/hdiskA1 and /dev/hdiskB1. They are in separate replicated failure groups, so any data stored on /dev/hdiskA1 is replicated on /dev/hdiskB1. The file descriptor quorum for **sql1ib_shared** is handled through /dev/hdiskC1.

5. Create NSDs for the disks to be used for the log file system.

At this point, storage replication is configured for **sql1ib_shared**, but it needs to be configured for the database and transaction logs. Next, create NSDs using the disks for logfs, ensuring they are assigned to the correct failure groups.

Create a file */tmp/nsdForLogfs1.txt*.

```

root@hostA1: /> cat /tmp/nsdForLogfs1.txt
/dev/hdiskA2:hostA1,hostA2,hostA3::dataAndMetadata:1
/dev/hdiskB2:hostB1,hostB2,hostB3::dataAndMetadata:2
/dev/hdiskC2:T::descOnly:3

```

```

root@hostA1: /> /usr/lpp/mmfs/bin/mmcrnsd -F /tmp/nsdForLogfs1.txt
mmcrnsd: Processing disk hdiskA2
mmcrnsd: Processing disk hdiskB2
mmcrnsd: Processing disk hdiskC2
mmcrnsd: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.

```

Verify that the NSDs have been created:

```

root@hostA1: /> /usr/lpp/mmfs/bin/mmlnsd -X

gpfs1004nsd 091A33584D675EDA /dev/hdiskA2 hdisk hostA1.torolab.ibm.com server node
gpfs1004nsd 091A33584D675EDA /dev/hdiskA2 hdisk hostA2.torolab.ibm.com server node
gpfs1004nsd 091A33584D675EDA /dev/hdiskA2 hdisk hostA3.torolab.ibm.com server node
gpfs1004nsd 091A33584D675EDA /dev/hdiskA2 hdisk hostB1.torolab.ibm.com

```



```

gpfs1004nsd 091A33584D675EDA /dev/hdiskA2 hdisk hostB2.torolab.ibm.com
gpfs1004nsd 091A33584D675EDA /dev/hdiskA2 hdisk hostB3.torolab.ibm.com
gpfs1005nsd 091A336D4D675EDC /dev/hdiskB2 hdisk hostA1.torolab.ibm.com
gpfs1005nsd 091A336D4D675EDC /dev/hdiskB2 hdisk hostA2.torolab.ibm.com
gpfs1005nsd 091A336D4D675EDC /dev/hdiskB2 hdisk hostA3.torolab.ibm.com
gpfs1005nsd 091A336D4D675EDC /dev/hdiskB2 hdisk hostB1.torolab.ibm.com server node
gpfs1005nsd 091A336D4D675EDC /dev/hdiskB2 hdisk hostB2.torolab.ibm.com server node
gpfs1005nsd 091A336D4D675EDC /dev/hdiskB2 hdisk hostB3.torolab.ibm.com server node
gpfs1006nsd 091A33434D675EE0 /dev/hdiskC2 hdisk T.torolab.ibm.com server node

```

6. Create the replicated logfs system.

In Step 5, GPFS rewrites /tmp/nsdForLogfs1.txt to include the nsd names, instead of the *hdisknames*. This is done by commenting out the entries made in /tmp/nsdForLogfs1.txt and adding entries required for creating the filesystem. Once the GPFS rewrites the file, it will read as follows:

```

root@hostA1:/> cat /tmp/nsdForLogfs1.txt
# /dev/hdiskA2:hostA1,hostA2,hostA3::dataAndMetadata:1
gpfs1004nsd:::dataAndMetadata:1::
# /dev/hdiskB2:hostB1,hostB2,hostB3::dataAndMetadata:2
gpfs1005nsd:::dataAndMetadata:2::
# /dev/hdiskC2:T::descOnly:3
gpfs1006nsd:::descOnly:3::

```

Create the **logfs** file system, containing 2 replicas, a disk block size of 1MB, maximum of 255 nodes, and mounted under **/logfs**:

```

root@hostA1:/> /usr/lpp/mmfs/bin/mmcfs logfs -F /tmp/nsdForLogfs1.txt -m 2 -M 2 -r 2 -R 2 -B 1M -n 255 -T /logfs

```

```

The following disks of logfs will be formatted on node hostB2:
gpfs1004nsd: size 438304768 KB
gpfs1005nsd: size 34603008 KB
gpfs1006nsd: size 57344 KB
Formatting file system ...
Disks up to size 6.7 TB can be added to storage pool 'system'.
Creating Inode File
Creating Allocation Maps
Clearing Inode Allocation Map
Clearing Block Allocation Map
Formatting Allocation Map for storage pool 'system'
Completed creation of file system /dev/logfs.
mmcfs: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.

```

Verify that the file system has been created with the disks in the proper failure groups:

```

root@hostA1:/> /usr/lpp/mmfs/bin/mmlsdisk logfs -L

```

```

disk driver sector failure holds holds storage
name type size group metadata data status availability disk ID pool remarks
-----
gpfs1004nsd nsd 512 1 yes yes ready up 1 system desc
gpfs1005nsd nsd 512 2 yes yes ready up 2 system desc
gpfs1006nsd nsd 512 3 no no ready up 3 system desc
Number of quorum disks: 3
Read quorum value: 2
Write quorum value: 2

```

7. Create NSDs for **datafs**, and create the **datafs** file system.

Create the NSDs for the database container file systems. For this file system, use 5 disks at each of the two main sites, as well as one disk for file system quorum from host T:

```

root@hostA1:/> cat /tmp/nsdForDatafs.txt
/dev/hdiskA3:hostA1,hostA2,hostA3::dataAndMetadata:1
/dev/hdiskA4:hostA1,hostA2,hostA3::dataAndMetadata:1
/dev/hdiskA5:hostA1,hostA2,hostA3::dataAndMetadata:1
/dev/hdiskA6:hostA1,hostA2,hostA3::dataAndMetadata:1
/dev/hdiskA7:hostA1,hostA2,hostA3::dataAndMetadata:1
/dev/hdiskB3:hostB1,hostB2,hostB3::dataAndMetadata:2
/dev/hdiskB4:hostB1,hostB2,hostB3::dataAndMetadata:2
/dev/hdiskB5:hostB1,hostB2,hostB3::dataAndMetadata:2
/dev/hdiskB6:hostB1,hostB2,hostB3::dataAndMetadata:2
/dev/hdiskB7:hostB1,hostB2,hostB3::dataAndMetadata:2
/dev/hdiskC3:T::descOnly:3

```

```

root@hostA1:/> /usr/lpp/mmfs/bin/mmcrnsd -F /tmp/nsdForDatafs.txt
mmcrnsd: Processing disk hdiskA3
mmcrnsd: Processing disk hdiskA4
mmcrnsd: Processing disk hdiskA5
mmcrnsd: Processing disk hdiskA6
mmcrnsd: Processing disk hdiskA7
mmcrnsd: Processing disk hdiskB3
mmcrnsd: Processing disk hdiskB4
mmcrnsd: Processing disk hdiskB5
mmcrnsd: Processing disk hdiskB6
mmcrnsd: Processing disk hdiskB7
mmcrnsd: Processing disk hdiskC3
mmcrnsd: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.

```

Create the **datafs** file system, containing two replicas, a disk block size of 1 MB, maximum of 255 nodes, and mounted under **/datafs**:

```

root@hostA1:/> /usr/lpp/mmfs/bin/mmcrfs datafs -F /tmp/nsdForDatafs.txt -m 2 -M 2 -r 2 -R 2 -B 1M -n
255 -T /datafs

```

```

The following disks of datafs will be formatted on node hostA3:
gpfs1016nsd: size 438304768 KB
gpfs1017nsd: size 438304768 KB
gpfs1018nsd: size 438304768 KB
gpfs1019nsd: size 1462220800 KB
gpfs1020nsd: size 1462220800 KB
gpfs1021nsd: size 157286400 KB
gpfs1022nsd: size 157286400 KB
gpfs1023nsd: size 157286400 KB
gpfs1024nsd: size 157286400 KB
gpfs1025nsd: size 157286400 KB
gpfs1026nsd: size 57344 KB
Formatting file system ...
Disks up to size 18 TB can be added to storage pool 'system'.
Creating Inode File
Creating Allocation Maps
Clearing Inode Allocation Map
Clearing Block Allocation Map
Formatting Allocation Map for storage pool 'system'
Completed creation of file system /dev/datafs.
mmcrfs: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.

```

8. Mount log file systems and data file systems.

```

root@hostA1:/> /usr/lpp/mmfs/bin/mmlsmount logfs
File system logfs is not mounted.

root@hostA1:/> /home/db2inst1/sqllib/bin/db2cluster -cfs -mount -filesystem logfs
File system 'logfs' was successfully mounted.

root@hostA1:/> /usr/lpp/mmfs/bin/mmlsmount logfs
File system logfs is mounted on 7 nodes.

root@hostA1:/> /home/db2inst1/sqllib/bin/db2cluster -cfs -mount -filesystem datafs
File system 'datafs' was successfully mounted.

```

9. Complete the affinization of Reads.

As root, complete the affinization of reads to local hosts by issuing the following command:

```

root@hostA1:/> mmchconfig readReplicaPolicy=local
mmchconfig: Command successfully completed
mmchconfig: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.

```

Creating the database in a GDPC environment

In a geographically dispersed DB2 pureScale cluster (GDPC) environment, create a database table with the datafs file system, with database logs stored in the logfs file system.

Procedure

Create the database.

As root, change the permissions on the file systems to allow the instance user to write to them:

```
root@hostA1:/> mkdir /logfs/db2inst1
root@hostA1:/> chown db2inst1 /logfs/db2inst1
root@hostA1:/> mkdir /datafs/db2inst1
root@hostA1:/> chown db2inst1 /datafs/db2inst1
```

```
db2inst1@hostA1:/home/db2inst1> db2start
02/25/2011 04:00:44 0 0 SQL1063N DB2START processing was successful.
02/25/2011 04:00:46 3 0 SQL1063N DB2START processing was successful.
02/25/2011 04:00:47 2 0 SQL1063N DB2START processing was successful.
02/25/2011 04:00:47 1 0 SQL1063N DB2START processing was successful.
SQL1063N DB2START processing was successful.
```

```
db2inst1@hostA1:/home/db2inst1> db2 create database MYDB on /datafs/db2inst1 dbpath
on /logfs/db2inst1
```

DB20000I The CREATE DATABASE command completed successfully.

Configuring the cluster for high availability in a GDPC environment

The configuration procedure detailed in this topic is specific to the geographically dispersed DB2 pureScale cluster (GDPC).

Procedure

1. Update storage failure time-outs.

Ensure that in the case of storage controller or site failure, an error is returned quickly to GPFS by setting the relevant device driver parameters. Note that the relevant parameters will differ for different device drivers. Check storage controller documentation or consult a storage expert on site to ensure that errors are returned within 20 seconds.

For DS3K/DS4K using the default AIX PCM, the updates are as follows:

```
chdev -l hdiskX -a 'cntl_delay_time=20 cntl_hcheck_int=2' -P
```

repeat for every hdiskx used by pureScale

```
chdev -l fscsiY -a dyntrk=yes -a fc_err_recov=fast_fail -P
```

repeat for every fscsiY adapter

reboot the host

repeat chdevs for every host in the cluster

Verify the attributes have been set correctly on every computer:

```
root> lsattr -El fscsi0
attach          switch          How this adapter is CONNECTED          False
dyntrk          yes             Dynamic Tracking of FC Devices          True
fc_err_recov    fast_fail      FC Fabric Event Error RECOVERY Policy  True

root> lsattr -El hdiskA1
PCM              PCM/friend/otherapdisk  Path Control Module          False
PR_key_value     none                 Persistent Reserve Key Value  True
Algorithm        fail_over            Algorithm                      True
autorecovery     no                   Path/Ownership Autorecovery  True
clr_q            no                   Device CLEARS its Queue on error True
cntl_delay_time  20                   Controller Delay Time         True
cntl_hcheck_int  2                     Controller Health Check Interval True
```

2. Update the resource time-outs.

Due to GPFS replication recovery requirements, recovery times for certain failures can be slightly longer in a geographically dispersed DB2 pureScale cluster (GDPC) environment than in a single-site pureScale environment. To account for this, some of the TSA resources need to have their timeout values adjusted. To adjust the time-outs, run the following commands once as root on any of the hosts in the cluster:

```
root> export CT_MANAGEMENT_SCOPE=2;
# Update 2 member-specific timeouts. For these, the resource
# names to update will look like db2_<instance>_<member_id>-rs.
# In this example we have members 0-4, and our instance name is
# db2inst1:
root> chrsrc -s "Name like 'db2_db2inst1_%-rs'" IBM.Application CleanupCommandTimeout=600
root> chrsrc -s "Name like 'db2_db2inst1_%-rs'" IBM.Application MonitorCommandTimeout=600

# In the next two commands, replace 'db2inst1' with your instance
# owning ID
root> chrsrc -s "Name like 'primary_db2inst1_900-rs'" IBM.Application CleanupCommandTimeout=600
root> chrsrc -s "Name like 'ca_db2inst1_0-rs'" IBM.Application CleanupCommandTimeout=600

# In the following commands, replace 'db2inst1' with your
# instance owning ID, and repeat for each host in your cluster,
# except the tiebreaker host T
root> chrsrc -s "Name like 'instancehost_db2inst1_hostA1'" IBM.Application MonitorCommandTimeout=600
root> chrsrc -s "Name like 'instancehost_db2inst1_hostA2'" IBM.Application MonitorCommandTimeout=600
root> chrsrc -s "Name like 'instancehost_db2inst1_hostA3'" IBM.Application MonitorCommandTimeout=600
root> chrsrc -s "Name like 'instancehost_db2inst1_hostB1'" IBM.Application MonitorCommandTimeout=600
root> chrsrc -s "Name like 'instancehost_db2inst1_hostB2'" IBM.Application MonitorCommandTimeout=600
root> chrsrc -s "Name like 'instancehost_db2inst1_hostB3'" IBM.Application MonitorCommandTimeout=600

# In the last two commands, replace 'db2inst1' with your instance
# owning ID, and 'hostA3' with the hostname of the first CF added
# to the cluster, and 'hostB3' with the hostname of the second
# CF added to the cluster.
root> chrsrc -s "Name like 'cacontrol_db2inst1_128_hostA3'" IBM.Application MonitorCommandTimeout=600
root> chrsrc -s "Name like 'cacontrol_db2inst1_129_hostB3'" IBM.Application MonitorCommandTimeout=600
```

To show the updated time-outs, run the following command as root:

```
lsrsrc -t IBM.Application Name MonitorCommandTimeout CleanupCommandTimeout
```

3. Update the condition/response scripts.

The installation of DB2 creates condition/response scripts that monitor the state of the system and react accordingly. In a single-site pureScale environment GPFS uses the Ethernet interfaces to communicate among its members and that Ethernet adapter is monitored by the condition/response scripts. In the GDPC environment, GPFS uses the InfiniBand (IB) or Remote Direct Memory Access (RDMA) over Converged Ethernet (RoCE) subnet for communication. Accordingly, the condition/response scripts need to be changed to reflect that.

List out the condition/response scripts:

```
root@hostA1:/> /home/db2inst1/sqllib/bin/db2cluster -cfs -list -network_resiliency -resources
```

For every host a condition will be listed that looks as follows:

```
condition 6:
    Name = "condrespV10_hostA1_condition_en2"
    Node = "hostA1.torolab.ibm.com"
    MonitorStatus = "Monitored"
    ResourceClass = "IBM.NetworkInterface"
    EventExpression = "OpState != 1"
    EventDescription = "Adapter is not online"
    RearmExpression = "OpState = 1"
    RearmDescription = "Adapter is online"
    SelectionString = "IPAddress == '9.26.82.X'"
    Severity = "c"
    NodeNames = {}
    MgtScope = ""
    Toggle = "Yes"
    EventBatchingInterval = 0
    EventBatchingMaxEvents = 0
    BatchedEventRetentionPeriod = 0
    BatchedEventMaxTotalSize = 0
    RecordAuditLog = "ALL"
```

The selection string will show the Ethernet address for the host. For every host except for the tiebreaker, the selection string needs to be updated to reflect the IB or RoCE address used on that host. In this example, the IB address on hostA1 is 10.1.1.1, so the command to update it is:

```
root@hostA1:/> chcondition -U condrespV10_hostA1_condition_en2
root@hostA1:/> chcondition -s "IPAddress == '10.1.1.1'" condrespV10_pscalep74008a_condition_en2
root@hostA1:/> chcondition -L condrespV10_hostA1_condition_en2
```

Repeat for every host except for the tiebreaker. Note the **chcondition** commands must be run on the host whose condition is being changed. An attempt to change the condition for hostA2 on hostA1 will result in a condition not found error.

Chapter 21. Removing DB2 data server and the DB2 pureScale Feature

Perform the steps in this task to remove the DB2 data server with the IBM DB2 pureScale Feature from your system.

Procedure

1. Back up your data.
2. Stop the DB2 pureScale instance by running the `db2stop` command from one of the hosts.
3. Drop the DB2 pureScale instance on all the hosts by running the following command from one of the hosts:

```
db2idrop -g instance_name
```

where the **-g** parameter specifies that the DB2 pureScale instance will be dropped globally (on all DB2 members and cluster caching facilities). You can find the **db2idrop** command in the `<DB2DIR>/instance` directory

4. Manually remove the IBM General Parallel File System (GPFS) cluster and file system. The **db2idrop -g** does not delete the GPFS cluster on the host from which you ran the command and does not remove file systems to protect your data.
5. Remove the DB2 data server with the DB2 pureScale Feature installation by running the following command:

```
db2_deinstall -a
```

Dropping a member or cluster caching facility

This task details the steps necessary to drop a member or cluster caching facility.

About this task

In a recoverable database, if you make any change to member topology, for example adding or dropping a member, you must take an offline backup before you can access the database. Otherwise, if you attempt to access the database before taking an offline backup, the database is placed in a backup pending state.

In a non-recoverable database, if you make any change to the member topology, you do not need to take an offline backup before you can access the database. However, the database would only be recoverable to the point in time when you took the last backup image, and using the same topology as when the image was produced. If you do not take an offline back and attempt to access the database, the database is not placed in a backup pending state.

You can add multiple members or drop multiple members without having to take a backup after each change. For example, if you add three members, you only have to take a backup after you have completed all add operations. However, if you add two members and then drop a member, or if you drop two members and then add a member, you must take a backup before you can perform any additional member topology changes.

Restrictions

The **db2iupdt -drop** command does not drop the last cluster caching facility or the last member, in the DB2 pureScale instance. To drop the last member or cluster caching facility in the DB2 pureScale instance, see the “Removing DB2 Enterprise Server Edition with the DB2 pureScale Feature” topic.

You must run the **db2iupdt -drop** command from a host that will still belong to the instance after you have dropped the cluster caching facility or member.

Procedure

1. Log in to the host that will still belong to the instance after you dropped the cluster caching facility or member.
2. Stop the DB2 pureScale instance on all hosts by using the **db2stop** command.
3. To remove a DB2 member:

```
db2iupdt -drop -m hostname instance_name
```

To remove a cluster caching facility:

```
db2iupdt -drop -cf hostname instance_name
```

4. Remove the DB2 pureScale Feature installation on the host by running the following command:

```
db2_deinstall -a
```

When the DB2 installer removes the last DB2 installation, it also automatically removes DB2 cluster services.

Example

For example, if you want to drop a member from a host called test1 and an instance called db2sdin1, run the following command:

```
db2iupdt -drop -m test1 db2sdin1
```

Then, to remove the DB2 installation from the test1 host, run the following command from the test1 host:

```
DB2DIR/install/db2_deinstall -a
```

Manually cleaning a DB2 managed clustered file system

This topic guides you through the required steps to manually clean a DB2 managed clustered file system.

About this task

Using the **db2idrop -g** command to remove the IBM DB2 pureScale Feature from your environment removes the GPFS cluster on all hosts except for the host on which the **db2idrop** command was run. Use this procedure to remove the GPFS file system and cluster on the remaining host. All data on the GPFS file system will be lost. After the **db2idrop** command has completed, the GPFS cluster will be left over on installation-initiating host (IIH) only. Manual clean up is only required on the host acting as the IIH.

Procedure

1. List existing GPFS file systems using the following command:

```
DB2DIR/bin/db2cluster -cfs -list -filesystem
```

where *DB2DIR* represents the installation location of your DB2 copy.

The output of this command should be similar to the following:

```
FILE SYSTEM NAME MOUNT_POINT
-----
db2fs1    /db2sd_20091027220651
```

2. Stop the entire GPFS cluster:

```
db2cluster -cfs -stop -all
```

3. Set the GPFS quorum type from tiebreaker to majority:

```
db2cluster -cfs -set -tiebreaker -majority
```

4. Start the GPFS cluster:

```
db2cluster -cfs -start -all
```

5. To ensure there is no data on the file system before deleting it, mount the file system:

```
db2cluster -cfs -mount -filesystem db2fs1
```

6. Delete the GPFS file system:

```
db2cluster -cfs -delete -filesystem db2fs1
```

The output of this command should be similar to the following:

```
The file system 'db2fs1' has been successfully deleted.
All cluster configurations have been completed successfully.
```

7. List the GPFS domain name:

```
db2cluster -cfs -list -domain
```

The output of this command should be similar to the following:

```
Domain Name: db2cluster_20091027220622.ca.ibm.com
```

8. Stop the GPFS cluster:

```
db2cluster -cfs -stop -all
```

9. Delete the GPFS cluster:

```
db2cluster -cfs -delete -domain db2cluster_20091027220622.ca.ibm.com
```

The output of this command should be similar to the following:

```
Deleting the domain db2cluster_20091027220622.in.ibm.com from
the cluster was successful.
```

10. After removing GPFS cluster and file systems, delete the GPFS_CLUSTER and DEFAULT_INSTPROF variable records in the Global Registry.

```
db2greg -delvarrec service=GPFS_CLUSTER,variable=NAME,installpath=-
db2greg -delvarrec service=DEFAULT_INSTPROF,variable=DEFAULT,installpath=-
```

Part 4. Installing the DB2 Partitioned Database Environment

To install the DB2 partitioned database environment, you need to set up a partitioned database environment and verify the installation.

After you complete the installation, take the necessary steps to ensure your environment is configured correctly.

Chapter 22. Setting up a partitioned database environment

This topic describes how to set up a partitioned database environment. You will use the DB2 Setup wizard to install your instance-owning database server and to create the response files that will in turn be used to create your participating database servers.

Before you begin

Note: A partitioned database environment is not supported in non-root installations.

- Ensure that you have the InfoSphere Warehouse Activation CD license key that will need to be copied over to all participating computers.
- The same number of consecutive ports must be free on each computer that is to participate in the partitioned database environment. For example, if the partitioned database environment will be comprised of four computers, then each of the four computers must have the same four consecutive ports free. During instance creation, a number of ports equal to the number of logical partitions on the current server will be reserved in the `/etc/services` on Linux and UNIX and in the `%SystemRoot%\system32\drivers\etc\services` on Windows. These ports will be used by the Fast Communication Manager. The reserved ports will be in the following format:

```
DB2_InstanceName
DB2_InstanceName_1
DB2_InstanceName_2
DB2_InstanceName_END
```

The only mandatory entries are the beginning (DB2_InstanceName) and ending (DB2_InstanceName_END) ports. The other entries are reserved in the services file so that other applications do not use these ports

- To support multiple participating DB2 database servers, the computer on which you want to install DB2 must belong to an accessible domain. However, you can add local partitions to this computer even though the computer doesn't belong to a domain.
- On Linux and UNIX systems, a remote shell utility is required for partitioned database systems. DB2 database systems support the following remote shell utilities:
 - rsh
 - ssh

By default, DB2 database systems use rsh when executing commands on remote DB2 nodes, for example, when starting a remote DB2 database partition. To use the DB2 default, the rsh-server package must be installed. For more information, see “Security considerations when installing and using the DB2 database manager” in *Database Security Guide*.

If you choose to use the rsh remote shell utility, inetd (or xinetd) must be installed and running as well. If you choose to use the ssh remote shell utility, you need to set the **DB2RSHCMD** registry variable immediately after the DB2 installation is complete. If this registry variable is not set, rsh is used.

- On Linux and UNIX operating systems, ensure the hosts file under the etc directory does not contain an entry for “127.0.0.2” if that IP address maps to the fully qualified hostname of the machine.

About this task

A database partition is part of a database that consists of its own data, indexes, configuration files, and transaction logs. A partitioned database is a database with two or more partitions.

Procedure

To set up a partitioned database environment:

1. Install your instance-owning database server using the DB2 Setup wizard. For detailed instructions, see the appropriate “Installing DB2 servers” topic for your platform.
 - On the **Select installation, response files creation, or both** window, ensure that you select the **Save my installation settings in a response files** option. After the installation has completed, two files will be copied to the directory specified in the DB2 Setup wizard: `PROD_ESE.rsp` and `PROD_ESE_addpart.rsp`. The `PROD_ESE.rsp` file is the response file for instance-owning database servers. The `PROD_ESE_addpart.rsp` file is the response file for participating database servers.
 - On the **Set up partitioning options for the DB2 instance** window, ensure that you select **Multiple partition instance**, and enter the maximum number of logical partitions.
2. Make the DB2 install image available to all participating computers in the partitioned database environment.
3. Distribute the participating database servers response file (`PROD_ESE_addpart.rsp`).
4. Install a DB2 database server on each of the participating computers using the **db2setup** command on Linux and UNIX, or the **setup** command on Windows:

Linux and UNIX

Go to the directory where the DB2 database product code is available and run:

```
./db2setup -r /responsefile_directory/response_file_name
```

Windows

```
setup -u x:\responsefile_directory\response_file_name
```

For example, here is the command using the `PROD_ESE_addpart.rsp` as the response file:

Linux and UNIX

Go to the directory where the DB2 database product code is available and run:

```
./db2setup -r /db2home/PROD_ESE_addpart.rsp
```

where `/db2home` is the directory where you have copied the response file.

Windows

```
setup -u c:\resp_files\PROD_ESE_addpart.rsp
```

where `c:\resp_files\` is the directory where you have copied the response file.

5. (Linux and UNIX only) Configure the `db2nodes.cfg` file. The DB2 installation only reserves the maximum number of logical partitions you want to use for

the current computer, but does not configure the `db2nodes.cfg` file. If you do not configure the `db2nodes.cfg` file, the instance is still a single partitioned instance.

6. Update the services file on the participating servers to define the corresponding FCM port for the DB2 instance. The services file is in the following location:
 - `/etc/services` on Linux and UNIX
 - `%SystemRoot%\system32\drivers\etc\services` on Windows
7. For partitioned database environments on Windows 2000 or later, start the DB2 Remote Command Service security feature to protect your data and resources. To be fully secure, start either the computer (if the service is running under the context of the LocalSystem account) or a user for delegation (if the service is being run under the logon context of a user).

To start the DB2 Remote Command Service security feature:

 - a. Open the Active Directory Users and Computers window on the domain controller, click **Start** and select **Programs > Administrative tools > Active Directory Users and Computers**
 - b. In the right window panel, right-click the computer or user to start, select **Properties**
 - c. Click the **General** tab and select the **Trust computer for delegation** check box. For user setting, click the **Account** tab and select the **Account is trusted for delegation** check box in the **Account option** group. Ensure that the **Account is sensitive and cannot be delegated** box has not been checked.
 - d. Click **OK** to start the computer or user for delegation.

Repeat these steps for each computer or user that needs to be started. You must restart your computer for the security change to take effect.

Chapter 23. Preparing to install the DB2 Database Partitioning Feature

To prepare your environment for a partitioned DB2 server you might need to update your settings or verify your port range availability on participating computers.

For Linux and UNIX, you might need to verify that the Network File System (NFS) is running on each computer participating in your partitioned database system.

For AIX, HP-UX, Linux and Solaris, you can create group and user IDs in a partitioned environment.

Preparing the environment for a partitioned DB2 server (Windows)

This topic describes the steps required to prepare your Windows environment for a partitioned installation of the DB2 database product.

Before you begin

When you add a new machine as a partition in a partitioned database environment, the new machine must:

- Have the same operating system version as the instance owning machine.
- Have the same CPU architecture (x32 bit or x64 bit) as the instance owning machine.

If the new machine does not meet these requirements, adding the partition might fail.

Procedure

To prepare your Windows environment for installation:

1. Ensure that the primary computer and participating computers belong to the same Windows domain. Check the domain to which the computer belongs by using the **System Properties** dialog, accessible through the Control Panel.
2. Ensure that time and date settings on the primary computer and participating computers are consistent. To be considered consistent, the difference in GMT time between all computers must be no greater than one hour.

System date and time can be modified using the **Date/Time Properties** dialog, accessible through the Control Panel. You can use the **max_time_diff** configuration parameter to change this restriction. The default is **max_time_diff = 60**, which allows a difference of less than 60 minutes.

3. Ensure that each computer object that participates in the partitioned database environment has the "Trust computer for delegation" privilege flagged. You can verify that the "Trust computer for delegation" check box on the **General** tab of each computer's account **Properties** dialog box in the Active Directory Users and Computers console is checked.
4. Ensure that all participating computers can communicate with each other using TCP/IP:
 - a. On one participating computer, enter the **hostname** command, which will return the hostname of the computer.

- b. On another participating computer, enter the following command:

```
ping hostname
```

where *hostname* represents the hostname of the primary computer. If the test is successful, you will receive output similar to the following:

Pinging ServerA.ibm.com [9.21.27.230] with 32 bytes of data:

```
Reply from 9.21.27.230: bytes=32 time<10ms TTL=128
```

```
Reply from 9.21.27.230: bytes=32 time<10ms TTL=128
```

```
Reply from 9.21.27.230: bytes=32 time<10ms TTL=128
```

Repeat these steps until you are sure that all participating computers can communicate with each other using TCP/IP. Each computer must have a static IP address.

If you are planning to use multiple network adapters, you can specify which adapter to use to communicate between database partition servers. Use the **db2nchg** command to specify the netname field in the `db2nodes.cfg` file after the installation is complete.

5. During the installation you will be asked to provide a DB2 Administration Server user account. This is a local or domain user account that will be used by the DB2 Administration Server (DAS). The DAS is an administration service used to support the GUI tools and assist with administration tasks. You can define a user now or have the DB2 Setup wizard create one for you. If you want to create a new domain user using the DB2 Setup wizard, the account used to perform the installation must have authority to create domain users.
6. On the primary computer, where you will install the instance-owning partition, you must have a domain user account that belongs to the local *Administrators* group. You will log on as this user when you install DB2 database products. You must add the same user account to the local *Administrators* group on each participating computer. This user must have the *Act as part of the operating system* user right.
7. Ensure that all computers in the instance have the database directory on the same local drive letter. You can check this condition by running the **GET DATABASE CONFIGURATION** command and verifying the value of the **dftdbpath** DBM configuration parameter.
8. During the installation you will be asked to provide a domain user account to be associated with the DB2 instance. Every DB2 instance has one user assigned. The DB2 database system logs on with this user name when the instance is started. You can define a user now, or you can have the DB2 Setup wizard create a new domain user for you.

When adding a new node to a partitioned environment the DB2 copy name must be the same on all computers.

If you want to create a new domain user using the DB2 Setup wizard, the account used to perform the installation must have authority to create domain users. The instance user domain account must belong to the local *Administrators* group on all the participating computers and will be granted the following user rights:

- Act as part of the operating system
- Create token object
- Lock pages in memory
- Log on as a service
- Increase quotas
- Replace a process level token

If extended security was selected, the account must also be a member of the DB2ADMNS group. The DB2ADMNS group already has these privileges so the privileges are already explicitly added to the account.

Updating environment settings for a partitioned DB2 installation (AIX)

This task describes the environment settings that you need to update on each computer that will participate in your partitioned database system.

Procedure

To update AIX environment settings:

1. Log on to the computer as a user with root user authority.
2. Set the AIX maxuproc (maximum number of processes per user) device attribute to 4096 by entering the following command:

```
chdev -l sys0 -a maxuproc='4096'
```

Note: A bosboot/reboot may be required to switch to the 64-bit kernel if a different image is being run.

3. Set the TCP/IP network parameters on all the workstations that are participating in your partitioned database system to the following values. These values are the minimum values for these parameters. If any of the network-related parameters are already set to a higher value, do not change it.

```
thewall      = 65536
sb_max       = 1310720
rfc1323      = 1
tcp_sendspace = 221184
tcp_recvspace = 221184
udp_sendspace = 65536
udp_recvspace = 65536
ipqmaxlen    = 250
somaxconn    = 1024
```

To list the current settings of all network-related parameters, enter the following command:

```
no -a | more
```

To set a parameter, enter the follow command:

```
no -o parameter_name=value
```

where:

- *parameter_name* represents the parameter you want to set.
- *value* represents the value that you want to set for this parameter.

For example, to set the tcp_sendspace parameter to 221184, enter the following command:

```
no -o tcp_sendspace=221184
```

4. If you are using a high speed interconnect, you must set the spoolsize and rpoolsize for css0 to the following values:

```
spoolsize    16777216
rpoolsize    16777216
```

To list the current settings of these parameters, enter the following command:

```
lsattr -l css0 -E
```

To set these parameters, enter the following commands:

```
/usr/lpp/ssp/css/chgcss -l css0 -a spoolsize=16777216
/usr/lpp/ssp/css/chgcss -l css0 -a rpoolsize=16777216
```

If you are not using the `/tftpboot/tuning.cst` file to tune your system, you can use the `DB2DIR/misc/rc.local.sample` sample script file, where `DB2DIR` is path where the DB2 database product has been installed, to update the network-related parameters after installation. To update the network-related parameters using the sample script file after installation, perform the following steps:

- a. Copy this script file to the `/etc` directory and make it executable by root by entering the following commands:

```
cp /usr/opt/db2_09_01/misc/rc.local.sample /etc/rc.local
chown root:sys /etc/rc.local
chmod 744 /etc/rc.local
```

- b. Review the `/etc/rc.local` file and update it if necessary.
- c. Add an entry to the `/etc/inittab` file so that the `/etc/rc.local` script is executed whenever the machine is rebooted. You can use the **mkinitab** command to add an entry to the `/etc/inittab` file. To add this entry, enter the following command:

```
mkinitab "rclocal:2:wait:/etc/rc.local > /dev/console 2>&1"
```

- d. Ensure that `/etc/rc.nfs` entry is included in the `/etc/inittab` file by entering the following command:

```
lsinitab rcnfs
```

- e. Update the network parameters without rebooting your system by entering the following command:

```
/etc/rc.local
```

5. Ensure that you have enough paging space for a partitioned installation of DB2 Enterprise Server Edition to run. If you do not have sufficient paging space, the operating system will kill the process that is using the most virtual memory (this is likely to be one of the DB2 processes). To check for available paging space, enter the following command:

```
lspas -a
```

This command will return output similar to the following:

Page Space	Physical Volume	Volume Group	Size	%Used	Active	Auto	Type
paging00	hdisk1	rootvg	60MB	19	yes	yes	lv
hd6	hdisk0	rootvg	60MB	21	yes	yes	lv
hd6	hdisk2	rootvg	64MB	21	yes	yes	lv

The paging space available should be equal to twice the amount of physical memory installed on your computer.

6. If you are creating a small to intermediate size partitioned database system, the number of network file system daemons (NFSDs) on the instance-owning computer should be close to:

of biod on a computer \times # of computers in the instance

Ideally, you should run 10 biod processes on every computer. According to the preceding formula, on a four computer system with 10 biod processes, you use 40 NFSDs.

If you are installing a larger system, you can have up to 120 NFSDs on the computer.

For additional information about NFS, refer to your NFS documentation.

Verifying port range availability on participating computers (Linux and UNIX)

This task describes the steps required to verify port range availability on participating computers. The port range is used by the Fast Communications Manager (FCM). FCM is a feature of DB2 that handles communications between database partition servers.

Before you begin

Verifying the port range availability on participating computers should be done after you install the instance-owning database partition server and before you install any participating database partition servers.

When you install the instance-owning database partition server on the primary computer, DB2 reserves a port range according to the specified number of logical database partition servers participating in partitioned database environment. The default range is four ports. For each server that participates in the partitioned database environment, you must manually configure the `/etc/services` file for the FCM ports. The range of the FCM ports depends on how many logical partitions you want to use on the participating computer. A minimum of two entries are required, `DB2_instance` and `DB2_instance_END`. Other requirements for the FCM ports specified on participating computers are:

- The starting port number must match the starting port number of the primary computer
- Subsequent ports must be sequentially numbered
- Specified port numbers must be free

To make changes to the `services` file, you require root user authority.

Procedure

To verify the port range availability on participating computers:

1. Open the `services` file located in the `/etc/services` directory.
2. Locate the ports reserved for the DB2 Fast Communications Manager (FCM). The entries should appear similar to the following example:

```
DB2_db2inst1      60000/tcp
DB2_db2inst1_1    60001/tcp
DB2_db2inst1_2    60002/tcp
DB2_db2inst1_END  60003/tcp
```

DB2 reserves the first four available ports after 60000.

3. On each participating computer, open the `services` file and verify that the ports reserved for DB2 FCM in the `services` file of the primary computer are not being used.
4. In the event that the required ports are in use on a participating computer, identify an available port range for all computers and update each service file, including the `services` file on the primary computer.

What to do next

After you install the instance-owning database partition server on the primary computer, you must install your DB2 database product on the participating database partition servers. You can use the response file generated for the

partitioning servers (default name is `db2ese_addpart.rsp`), you need to manually configure the `/etc/services` files for the FCM ports. The range of the FCM ports depend on how many logical partitions you want to use on the current machine. The minimum entries are for `DB2_` and `DB2__END` two entries with consecutive free port numbers. The FCM port numbers used on each participating machines must have the same starting port number, and subsequent ports must be sequentially numbered.

Creating a DB2 home file system for a partitioned database environment

To install DB2 Enterprise Server Edition and create a partitioned database environment, you must have a file system that is available to all of the machines.

This file system will be used as the instance home directory. It can also be used as a temporary location for the contents of the DB2 database product CD.

Creating a DB2 home file system for a partitioned database system (AIX)

As part of setting up your partitioned database system, you need to create a DB2 home file system. Then you must NFS export the home file system and mount it from each computer participating in the partitioned database system.

Before you begin

It is recommended that you create a home file system that is as large as the content on the DB2 database product DVD. You can use the following command to check the size, KB:

```
du -sk DVD_mounting_point
```

A DB2 instance will require at least 200 MB of space. If you do not have enough free space, you can mount the DB2 database product DVD from each participating computer as an alternative to copying the contents to disk.

You must have:

- root authority to create a file system
- Created a volume group where your file system is to be physically located.

Procedure

To create, NFS export, and NFS mount the DB2 home file system, perform the following steps:

1. Create the DB2 home file system.

Log on to the primary computer (ServerA) in your partitioned database system as a user with root authority and create a home file system for your partitioned database system called `/db2home`.

- a. Enter the **`smit jfs`** command.
- b. Click on the **Add a Journaled File System** icon.
- c. Click on the **Add a Standard Journaled File System** icon.
- d. Select an existing volume group from the **Volume Group Name** list where you want this file system to physically reside.

- e. Set the **SIZE** of file system (**SIZE of file system (in 512-byte blocks) (Num.)** field). This sizing is enumerated in 512-byte blocks, so if you only need to create a file system for the instance home directory, you can use 180 000, which is about 90 MB. If you need to copy the product DVD image over to run the installation, you can create it with a value of 2 000 000, which is about 1 GB.
 - f. Enter the mount point for this file system in the **MOUNT POINT** field. In this example, the mount point is /db2home.
 - g. Set the **Mount AUTOMATICALLY at system restart** field to yes.
The remaining fields can be left to the default settings.
 - h. Click **OK**.
2. Export the DB2 home file system.
NFS export the /db2home file system so that it is available to all of the computers that will participate in your partitioned database system.
 - a. Enter the **smit nfs** command.
 - b. Click on the **Network File System (NFS)** icon.
 - c. Click on the **Add a Directory to Exports List** icon.
 - d. Enter the path name and directory to export (for example, /db2home) in the **PATHNAME of directory to export** field.
 - e. Enter the name of each workstation that will participate in your partitioned database system in the **HOSTS allowed root access** field. Use a comma (,) as the delimiter between each name. For example, ServerA, ServerB, ServerC. If you are using a high speed interconnect, it is recommended that you specify the high speed interconnect names for each workstation in this field as well. The remaining fields can be left to the default settings.
 - f. Click **OK**.
 3. Log out.
 4. Mount the DB2 home file system from each participating computer.
Log on to *each* participating computer (ServerB, ServerC, ServerD) and NFS mount the file system that you exported by performing the following steps:
 - a. Enter the **smit nfs** command.
 - b. Click on the **Network File System (NFS)** icon.
 - c. Click on the **Add a File System for Mounting** icon.
 - d. Enter the path name of the mount point in the **PATHNAME of the mount point (Path)** field.
The path name of the mount point is where you should create the DB2 home directory. For this example, use /db2home.
 - e. Enter the path name of the remote directory in the **PATHNAME of the remote directory** field.
For this example, you should enter the same value that you entered in the **PATHNAME of the mount point (Path)** field.
 - f. Enter the *hostname* of the machine where you exported the file system in the **HOST where the remote directory resides** field.
This value is the hostname of the machine where the file system that you are mounting was created.
To improve performance, you may want to NFS mount the file system that you created over a high speed interconnect. If you want to mount this file system using a high speed interconnect, you must enter its name in the **HOST where remote directory resides** field.

You should be aware that if the high speed interconnect ever becomes unavailable for some reason, every workstation that participates in your partitioned database system will lose access to the DB2 home directory.

- g. Set the **MOUNT now, add entry to /etc/filesystems or both?** field to both.
- h. Set the **/etc/filesystems entry will mount the directory on system RESTART** field to yes.
- i. Set the **MODE for this NFS file system** field to read-write.
- j. Set the **Mount file system soft or hard** field to hard.

A soft mount means that the computer *will not* try for an infinite period of time to remotely mount the directory. A hard mount means that your machine will infinitely try to mount the directory. This could cause problems in the event of a system crash. It is recommended that you set this field to hard.

The remaining fields can be left to the default settings.

- k. Ensure that this file system is mounted with the **Allow execution of SUID and sgid programs in this file system?** field set to Yes. This is the default setting.
- l. Click **OK**.
- m. Log out.

Creating a DB2 home file system for a partitioned database system (HP-UX)

As part of setting up your partitioned database system, you need to create a DB2 home file system. Then you must NFS export the home file system and mount it from each computer participating in the partitioned database system.

Before you begin

It is recommended that you create a home file system that is as large as the content on the DB2 database product DVD. You can use the following command to check the size, shown in KB:

```
du -sk DVD_mounting_point
```

A DB2 instance will require at least 50 MB of space. If you do not have enough free space, you can mount the DB2 database product DVD from each participating computer as an alternative to copying the contents to disk.

You must have root authority to create a file system.

Procedure

To create, NFS export, and NFS mount the DB2 home file system, perform the following steps:

1. Create the DB2 home file system.
 - Manually:
 - a. Select a disk partition or logical volume and use a utility like **newfs** to create this file system. For more information, enter the **man newfs** command.
 - b. Mount this file system locally and add an entry to the **/etc/fstab** file so that this file system is mounted each time the system is restarted.
 - Using SAM:

- a. Enter the **sam** command.
 - b. Click on the **Disks and File Systems** icon.
 - c. Click on the **File Systems** icon.
 - d. Select **Action > Add Local File systems**.
 - e. You can choose either to use or not to use a Logical Volume Manager. It is recommended to use a Logical Volume Manager.
2. Export the DB2 home file system.

If you are installing DB2 Enterprise Server Edition on a cluster of HP-UX systems, you can add an entry to the `/etc/exports` file to export this file system via NFS, or use SAM.

To export the file system using SAM:

- a. Enter the **sam** command.
 - b. Click on the **Networking and Communications** icon.
 - c. Click on the **Networked File Systems** icon.
 - d. Click on the **Exported Local File Systems** icon.
 - e. Click the **Action** menu and select **Add Exported File System**
 - f. Enter the path name and directory to export (for example, `/db2home`) in the **Local Directory Name** field.
 - g. Click the **User Access** button and add read-write access for the other computers in the instance in the window that appears.
 - h. Click the **Root User Access** button and add access for the other computers in the instance in the window that appears.
 - i. Click **OK**.
3. Log out.
4. Mount the DB2 home file system from each participating computer.

After you have exported this file system, you must mount this file system on each of the participating computers.

On each participating computer:

- a. Enter the **sam** command.
 - b. Click on the **Networking and Communications** icon.
 - c. Click on the **Networked File Systems** icon.
 - d. Click on the **Mounted Remote File Systems** icon.
 - e. Click the **Action** menu and select **Add Remote File System Using NFS**
 - f. Enter the mount point of the file system to mount (for example, `/db2home`) in the **Local Directory Name** field.
 - g. Enter the name of the remote server (for example, `ServerA`) in the **Remote Server Name** field.
 - h. Enter the path name and directory of the remote directory (for example, `/db2home`) in the **Remote Directory Name** field.
 - i. Set the **Mount At System Boot** option on.
 - j. Click the **NFS Mount Options** button and set **hard** mount type and the **Allow SetUID Execution** option on.

A soft mount means that the computer *will not* try for an infinite period of time to remotely mount the directory. A hard mount means that your machine will infinitely try to mount the directory. This could cause problems in the event of a system crash. It is recommended that you set this field to hard.

The remaining fields can be left to the default settings.

- k. Click OK.
- l. Log out.

Creating a file system for a partitioned database system (Linux)

As part of setting up your partitioned database system on Linux operating systems, you need to create a DB2 home file system. Then you must NFS export the home file system and mount it from each computer participating in the partitioned database system.

About this task

You must have a file system that is available to all machines that will participate in your partitioned database system. This file system will be used as the instance home directory.

For configurations that use more than one machine for a single database instance, NFS (Network File System) is used to share this file system. Typically, one machine in a cluster is used to export the file system using NFS, and the remaining machines in the cluster mount the NFS file system from this machine. The machine that exports the file system has the file system mounted locally.

For more command information, see your Linux distribution documentation.

Procedure

To create, NFS export, and NFS mount the DB2 home file system, perform the following steps:

1. On one machine, select a disk partition or create one using **fdisk**.
2. Using a utility like **mkfs**, create a file system on this partition. The file system should be large enough to contain the necessary DB2 program files as well as enough space for your database needs.
3. Locally mount the file system you have just created and add an entry to the `/etc/fstab` file so that this file system is mounted each time the system is rebooted. For example:

```
/dev/hda1 /db2home ext3 defaults 1 2
```

4. To automatically export an NFS file system on Linux at boot time, add an entry to the `/etc/exports` file. Be sure to include all of the host names participating in the cluster as well as all of the names that a machine might be known as. Also, ensure that each machine in the cluster has root authority on the exported file system by using the "root" option.

The `/etc/exports` file is an ASCII file which contains the following type of information:

```
/db2home machine1_name(rw) machine2_name(rw)
```

To export the NFS directory, run

```
/usr/sbin/exportfs -r
```

5. On each of the remaining machines in the cluster, add an entry to the `/etc/fstab` file to NFS mount the file system automatically at boot time. As in the following example, when you specify the mount point options, ensure that the file system is mounted at boot time, is read-write, is mounted hard, includes the `bg` (background) option, and that **setuid** programs can be run properly.

```
fusion-en:/db2home /db2home nfs rw,timeo=7,  
hard,intr,bg,suid,lock
```

where *fusion-en* represents the machine name.

6. NFS mount the exported file system on each of the remaining machines in the cluster. Enter the following command:

```
mount /db2home
```

If the **mount** command fails, use the **showmount** command to check the status of the NFS server. For example:

```
showmount -e fusion-en
```

where *fusion-en* represents the machine name.

This **showmount** command should list the file systems which are exported from the machine named *fusion-en*. If this command fails, the NFS server may not have been started. Run the following command as root on the NFS server to start the server manually:

```
/etc/rc.d/init.d/nfs restart
```

Assuming the present run level is 3, you can have this command run automatically at boot time by renaming *K20nfs* to *S20nfs* under the following directory: */etc/rc.d/rc3.d*.

Results

By performing these steps, you have completed the following tasks:

1. On a single computer in the partitioned database environment, you have created a file system to be used as the instance and home directory.
2. If you have a configuration that uses more than one machine for a single database instance, you have exported this file system using NFS.
3. You have mounted the exported file system on each participating computer.

Creating a DB2 home file system for a partitioned database system (Solaris)

As part of setting up your partitioned database system on Solaris Operating Environment, you need to create a DB2 home file system. Then you must NFS export the home file system and mount it from each computer participating in the partitioned database system.

Before you begin

It is recommended that you create a home file system that is as large as the content on the DB2 database product DVD. You can use the following command to check the size, shown in KB:

```
du -sk DVD mounting point
```

A DB2 instance will require at least 50 MB of space. If you do not have enough free space, you can mount the DB2 database product DVD from each participating computer as an alternative to copying the contents to disk.

You must have root authority to create a file system.

About this task

There are a number of ways to create a local file system on a Solaris Operating Environment system. If you want to use a product such as Veritas to create the file system, refer to the product's documentation.

Procedure

To create, NFS export, and NFS mount the DB2 home file system, perform the following steps:

1. Create the DB2 home file system.
 - a. On the primary computer (ServerA), select a disk partition or configure one using the **format** command.

When using the **format** command, ensure that the disk partitions being used do not overlap. Overlapping partitions can cause data corruption or file system failures. Ensure you have correctly entered the command, as mistakes can cause serious problems.
 - b. Using a utility like **newfs** or **mkfs**, create a file system on this partition.

The file system should be large enough to contain the necessary DB2 database files as well as other non-DB2 files. A minimum of 300 MB is recommended.
 - c. Locally mount the file system you have just created and add an entry to the `/etc/vfstab` file so that this file system is mounted each time the system is rebooted. For example:

```
/dev/dsk/clt0d2s2 /dev/rdisk/clt0d2s2 /db2home ufs 2 yes -
```

2. Export the DB2 home file system.
 - a. To automatically export an NFS file system on Solaris at boot time, add an entry to the `/etc/dfs/dfstab` file. Be sure to include all of the host names of the participating computers as well as all of the names that a given computer might be known as. Also, ensure that each computer has root authority on the exported file system by using the "root" option.

In the following example, an entry for a four computer partitioned database system is added to the `/etc/dfs/dfstab` file. The participating computers, ServerB, ServerC, and ServerD, are given permission to mount the file system `/db2home`, which will be used as the DB2 home file system.

```
share -F nfs -o \
rw=ServerB.torolab.ibm.com,\
root=ServerB.torolab.ibm.com \

rw=ServerC.torolab.ibm.com, \
root=ServerC.torolab.ibm.com\

rw=ServerD.torolab.ibm.com,\
root=ServerD.torolab.ibm.com \
-d "homes" /db2home
```

If a computer is known by more than one hostname, all aliases must be included in the `/etc/dfs/dfstab` file. For example, if ServerB was also known by the name ServerB-tokenring, the entry in the `/etc/dfs/dfstab` for ServerB would appear as follows:

```
rw=ServerB.torolab.ibm.com:ServerB-tokenring.torolab.ibm.com,\
root=ServerB.torolab.ibm.com:ServerB-tokenring.torolab.ibm.com \
```

- b. On each of the participating computers, add an entry to the `/etc/vfstab` file to NFS mount the file system automatically at boot time. As in the following example, when you specify the mount point options, ensure that

the file system is mounted at boot time, is read-write, is mounted hard, includes the **bg** (background) option, and that **suid** programs can be run properly:

```
ServerA:/db2home - /db2home nfs - yes rw,hard,intr,bg,suid
```

3. Mount the DB2 home file system from each participating computer.

On each of the participating computers in the partitioned database environment, enter the following commands:

```
mkdir /db2home
mount /db2home
```

If the **mount** command fails, use the **showmount** command to check the status of the NFS server. For example:

```
showmount -e ServerA
```

This **showmount** command should list the file systems that are exported from the computer named ServerA. If this command fails, the NFS server may not have been started. To start the server manually, run the following commands as root on the NFS server:

```
/usr/lib/nfs/mountd
/usr/lib/nfs/nfsd -a 16
```

These commands are run automatically at boot time if there are any entries in the `/etc/dfs/dfstab` file. After starting the NFS server, export the NFS file system again by running the following command:

```
sh /etc/dfs/dfstab
```

Results

By performing these steps, you have completed the following tasks:

1. On a single computer in the partitioned database environment, you have created a file system to be used as the instance and home directory.
2. You have exported this file system via NFS.
3. You have mounted the exported file system on each participating computer.

Verifying that NFS is running (Linux and UNIX)

Before setting up a database partitioned environment, you should verify that Network File System (NFS) is running on each computer that will participate in your partitioned database system.

Procedure

To verify that NFS is running on each computer:

- AIX operating systems:

Type the following command on each computer:

```
lssrc -g nfs
```

The Status field for NFS processes should indicate active. After you have verified that NFS is running on each system, you should check for the specific NFS processes required by DB2 database products. The required processes are:

```
rpc.lockd
rpc.statd
```

- HP-UX and Solaris operating systems:

Type the following command on each computer:

```
showmount -e hostname
```

Enter the **showmount** command without the *hostname* parameter to check the local system. If NFS is not active you will receive a message similar to the following:

```
showmount: ServerA: RPC: Program not registered
```

After you have verified that NFS is running on each system, you should check for the specific NFS processes required by DB2 database products:

```
rpc.lockd  
rpc.statd
```

You can use the following commands to check for these processes:

```
ps -ef | grep rpc.lockd  
ps -ef | grep rpc.statd
```

- Linux operating systems:

Type the following command on each computer:

```
showmount -e hostname
```

Enter the **showmount** command without the *hostname* parameter to check the local system.

If NFS is not active you will receive a message similar to the following:

```
showmount: ServerA: RPC: Program not registered
```

After you have verified that NFS is running on each system, you should check for the specific NFS processes required by DB2 database products. The required process is `rpc.statd`.

You can use the **ps -ef | grep rpc.statd** commands to check for this process.

If these processes are not running, consult your operating system documentation.

Creating groups and users in a partitioned database environment

To install the IBM DB2 Enterprise Server Edition software and create a partitioned database environment, you must create users and groups to operate the DB2 databases.

Creating required users for a DB2 server installation in a partitioned database environment (AIX)

Three users and groups are required to operate DB2 databases in partitioned database environments on AIX operating systems.

Before you begin

- You must have root user authority to create users and groups.
- If you manage users and groups with security software, additional steps might be required when defining DB2 users and groups.

About this task

The user and group names used in the following instructions are documented in the following table. You can specify your own user and group names if they adhere to your system naming rules and DB2 naming rules.

If you are planning to use the DB2 Setup wizard to install your DB2 database product, the DB2 Setup wizard will create these users for you.

Table 43. Required users and groups

Required user	User name	Group name
Instance owner	db2inst1	db2iadm1
Fenced user	db2fenc1	db2fadm1
DB2 administration server user	dasusr1	dasadm1

If the DB2 administration server user is an existing user, this user must exist on all the participating computers before the installation. If you use the DB2 Setup wizard to create a new user for the DB2 administration server on the instance-owning computer, then the new user is also created (if necessary) during the response file installations on the participating computers. If the user already exists on the participating computers, the user must have the same primary group.

Restrictions

The user names you create must conform to both your operating system's naming rules, and those of the DB2 database system.

Procedure

To create all three of these users, perform the following steps:

1. Log on to the primary computer.
2. Create a group for the instance owner (for example, db2iadm1), the group that will run UDFs or stored procedures (for example, db2fadm1), and the group that will own the DB2 administration server (for example, dasadm1) by entering the following commands:

```
mkgroup id=999 db2iadm1
mkgroup id=998 db2fadm1
mkgroup id=997 dasadm1
```

3. Create a user that belongs to each group that you created in the previous step using the following commands. The home directory for each user will be the DB2 home directory that you previously created and shared (db2home).

```
mkuser id=1004 pgrp=db2iadm1 groups=db2iadm1 home=/db2home/db2inst1
core=-1 data=491519 stack=32767 rss=-1 fsize=-1 db2inst1
mkuser id=1003 pgrp=db2fadm1 groups=db2fadm1 home=/db2home/db2fenc1
db2fenc1
mkuser id=1002 pgrp=dasadm1 groups=dasadm1 home=/home/dasusr1
dasusr1
```

4. Set an initial password for each user that you created by entering the following commands:

```
passwd db2inst1
passwd db2fenc1
passwd dasusr1
```

5. Log out.
6. Log on to the primary computer as each user that you created (db2inst1, db2fenc1, and dasusr1). You might be prompted to change each user's password because this is the first time that these users have logged onto the system.
7. Log out.
8. Create the exact same user and group accounts on each computer that will participate in your partitioned database environment.

Creating required users for a DB2 server installation in a partitioned database environment (HP-UX)

Three users and groups are required to operate DB2 databases in partitioned database environments on HP-UX operating systems.

Before you begin

- You must have root user authority to create users and groups.
- If you manage users and groups with security software, additional steps might be required when defining DB2 users and groups.

About this task

The user and group names used in the following instructions are documented in the following table. You can specify your own user and group names if they adhere to your system naming rules and DB2 naming rules.

If you are planning to use the DB2 Setup wizard to install your DB2 database product, the DB2 Setup wizard will create these users for you.

Table 44. Required users and groups

Required user	User name	Group name
Instance owner	db2inst1	db2iadm1
Fenced user	db2fenc1	db2fadm1
DB2 administration server user	dasusr1	dasadm1

If the DB2 administration server user is an existing user, this user must exist on all the participating computers before the installation. If you use the DB2 Setup wizard to create a new user for the DB2 administration server on the instance-owning computer, then the new user is also created (if necessary) during the response file installations on the participating computers. If the user already exists on the participating computers, the user must have the same primary group.

Restrictions

The user names you create must conform to both your operating system's naming rules, and those of the DB2 database system.

Procedure

To create all three of these users, perform the following steps:

1. Log on to the primary computer.
2. Create a group for the instance owner (for example, db2iadm1), the group that will run UDFs or stored procedures (for example, db2fadm1), and the group that will own the DB2 administration server (for example, dasadm1) by entering the following commands:

```
groupadd -g 999 db2iadm1
groupadd -g 998 db2fadm1
groupadd -g 997 dasadm1
```
3. Create a user that belongs to each group that you created in the previous step using the following commands. The home directory for each user will be the DB2 home directory that you previously created and shared (db2home).


```

useradd -u 1004 db2iadm1 -d /db2home/db2inst1 -m db2inst1
useradd -u 1003 db2fadm1 -d /db2home/db2fenc1 -m db2fenc1
useradd -u 1002 dasadm1 -d /home/dasusr1 -m dasusr1

```

4. Set an initial password for each user that you created by entering the following commands:


```

passwd db2inst1    passwd db2fenc1    passwd dasusr1

```
5. Log out.
6. Log on to the primary computer as each user that you created (db2inst1, db2fenc1, and dasusr1). You might be prompted to change each user's password because this is the first time that these users have logged onto the system.
7. Log out.
8. Create the exact same user and group accounts on each computer that will participate in your partitioned database environment.

Required users for a DB2 pureScale Feature installation (Linux)

Two users and groups are required to operate a DB2 databases environment on Linux operating systems.

Before you begin

- You must have root user authority to create users and groups.
- If you manage users and groups with security software, additional steps might be required when defining DB2 users and groups.

About this task

You need two users to create the DB2 pureScale instance:

- One user for the instance owner
- One user for the fenced user

You should use two different users with two different groups. Each of the two users should have the same UID, GID, group name, and the home directory on all the hosts. Ensure that if any of the users to be used exist on any of the hosts, they much have matching properties. It is not necessary to create these required users before you start the installation. You can create these users as you proceed through the panels of the DB2 Setup wizard or specify them in your response file. If existing users are used, they have to exist on all the hosts and meet the listed requirements.

The user and group names used in the following instructions are the defaults, and are documented in the following table. You can specify your own user and group names as long as they adhere to your system naming rules and DB2 naming rules.

Table 45. Default users and groups

Required user	User name	Group name
Instance owner	db2sdin1	db2iadm1
Fenced user	db2sdfe1	db2fadm1

The user and group names used in the following instructions are documented in the following table. You can specify your own user and group names if they adhere to your system naming rules and DB2 naming rules.

If you are planning to use the DB2 Setup wizard to install your DB2 database product, the DB2 Setup wizard will create these users for you.

Restrictions

The user names you create must conform to both your operating system's naming rules, and those of the DB2 database system.

The same user name that you will create on different hosts must have the same HOME directory. However, the user names must not already exist on any host. If existing user names are used, the user names must exist on all hosts with the same user ID (uid), group ID (gid), group name, and HOME directory.

Procedure

To create these users, perform the following steps:

1. Log on to a host.
2. Create a group for the instance owner (for example, db2iadm1) and a group that will run UDFs or stored procedures (for example, db2fadm1) by entering the following commands:

```
groupadd -g 999 db2iadm1
groupadd -g 998 db2fadm1
```

Ensure that the specific numbers you are using do not currently exist on any of the machines.

3. Create a user that belongs to each group that you created in the previous step using the following commands. The home directory for each user will be the DB2 home directory that you previously created and shared (db2home).

```
useradd -u 1004 -g db2iadm1 -m -d /db2home/db2inst1 db2inst1
useradd -u 1003 -g db2fadm1 -m -d /db2home/db2fenc1 db2fenc1
```

4. Set an initial password for each user that you created by entering the following commands:

```
passwd db2inst1    passwd db2fenc1
```

5. Log out.
6. Log on to the primary computer as each user that you created (db2inst1 and db2fenc1). You might be prompted to change each user's password because this is the first time that these users have logged onto the system.
7. Log out.
8. Create the exact same user and group accounts on each computer that will participate in your database environment.

Creating required users for a DB2 server installation in a partitioned database environment (Solaris Operating System)

Three users and groups are required to operate DB2 databases in partitioned database environments on Solaris operating systems.

Before you begin

- You must have root user authority to create users and groups.
- If you manage users and groups with security software, additional steps might be required when defining DB2 users and groups.

About this task

The user and group names used in the following instructions are documented in the following table. You can specify your own user and group names if they adhere to your system naming rules and DB2 naming rules.

If you are planning to use the DB2 Setup wizard to install your DB2 database product, the DB2 Setup wizard will create these users for you.

Table 46. Required users and groups

Required user	User name	Group name
Instance owner	db2inst1	db2iadm1
Fenced user	db2fenc1	db2fadm1
DB2 administration server user	dasusr1	dasadm1

If the DB2 administration server user is an existing user, this user must exist on all the participating computers before the installation. If you use the DB2 Setup wizard to create a new user for the DB2 administration server on the instance-owning computer, then the new user is also created (if necessary) during the response file installations on the participating computers. If the user already exists on the participating computers, the user must have the same primary group.

Restrictions

The user names you create must conform to both your operating system's naming rules, and those of the DB2 database system.

Procedure

To create all three of these users, perform the following steps:

1. Log on to the primary computer.
2. Create a group for the instance owner (for example, db2iadm1), the group that will run UDFs or stored procedures (for example, db2fadm1), and the group that will own the DB2 administration server (for example, dasadm1) by entering the following commands:

```
groupadd -g 999 db2iadm1
groupadd -g 998 db2fadm1
groupadd -g 997 dasadm1
```

3. Create a user for each group using the following commands. The home directory for each user will be the DB2 home directory that you previously created and shared (/db2home).

```
useradd -u 1004 -g db2iadm1 -d /db2home/db2inst1 -m db2inst1
useradd -u 1003 -g db2fadm1 -d /db2home/db2fenc1 -m db2fenc1
useradd -u 1002 -g dasadm1 -d /export/home/dasusr1 -m dasusr1
```

4. Set an initial password for each user that you created by entering the following commands:

```
passwd db2inst1    passwd db2fenc1    passwd dasusr1
```

5. Log out.
6. Log on to the primary computer as each user that you created (db2inst1, db2fenc1, and dasusr1). You might be prompted to change each user's password because this is the first time that these users have logged onto the system.

7. Log out.
8. Create the exact same user and group accounts on each computer that will participate in your partitioned database environment.

Chapter 24. Verifying the installation

You should verify that your DB2 database server installation was successful.

Verifying access to the registry on the instance-owning computer (Windows)

After installing DB2 database products in a partitioned database environment, verify that you have access to the registry on the instance-owning computer. This verification step must be run from all participating computers.

Before you begin

This step is only required if you are setting up a partitioned database environment.

You must have a DB2 database product installed.

Procedure

To verify access to the registry on the instance-owning computer:

1. From a command prompt, type the **regedit** command. The Registry Editor window opens.
2. Click the **Registry** menu item.
3. Select **Connect Network Registry**.
4. Type the name of the instance-owning computer in the **Computer name** field.

Results

If the information is returned, then you have successfully verified that you have access to the registry on the instance-owning computer.

Verifying a partitioned database server installation (Linux and UNIX)

Use the **db2va1** tool to verify the core functions of a DB2 copy by validating installation files, instances, database creation, connections to that database, and the state of partitioned database environments.

For details, see “Validating your DB2 copy”. The state of a partitioned database environment is only verified if there are at least 2 nodes. In addition, to verify that your DB2 database server installation was successful, you will create a sample database and run SQL commands to retrieve sample data and to verify that the data has been distributed to all participating database partition servers.

Before you begin

Before following these steps, make sure you have completed all of the installation steps.

Procedure

To create the SAMPLE database:

1. Log on to the primary computer (ServerA) as the instance-owning user. For this example, db2inst1 is the instance-owning user.
2. Enter the **db2samp1** command to create the SAMPLE database. By default, the sample database will be created in the instance-owner's home directory. In our example /db2home/db2inst1/ is the instance owner's home directory. The instance owner's home directory is the default database path.

This command can take a few minutes to process. There is no completion message; when the command prompt returns, the process is complete.

The SAMPLE database is automatically cataloged with the database alias SAMPLE when it is created.

3. Start the database manager by entering the **db2start** command.
4. Enter the following DB2 commands from a DB2 command window to connect to the SAMPLE database, retrieve a list of all the employees that work in department 20:

```
db2 connect to sample
db2 "select * from staff where dept = 20"
```

5. To verify that data has been distributed across database partition servers, enter the following commands from a DB2 command window:

```
db2 "select distinct dbpartitionnum(empno) from employee"
```

The output will list the database partitions used by the employee table. The specific output will depend on:

- The number of database partitions in the database
- The number of database partitions in the database partition group that is used by the table space where the employee table was created

What to do next

After you have verified the installation, you can remove the SAMPLE database to free up disk space. Enter the **db2 drop database sample** command to drop the SAMPLE database.

Verifying a partitioned database environment installation (Windows)

To verify that your DB2 database server installation was successful, you will create a sample database and run SQL commands to retrieve sample data and to verify that the data has been distributed to all participating database partition servers.

Before you begin

You have completed all of the installation steps.

Procedure

To create the SAMPLE database:

1. Log on to the primary computer (ServerA) as user with SYSADM authority.
2. Enter the **db2samp1** command to create the SAMPLE database.

This command can take a few minutes to process. When the command prompt returns, the process is complete.

The SAMPLE database is automatically cataloged with the database alias SAMPLE when it is created.

3. Start the database manager by entering the **db2start** command.

4. Enter the following DB2 commands from a DB2 command window to connect to the SAMPLE database, retrieve a list of all the employees that work in department 20:

```
db2 connect to sample
db2 "select * from staff where dept = 20"
```

5. To verify that data has been distributed across database partition servers, enter the following commands from a DB2 command window:

```
db2 "select distinct dbpartitionnum(empno) from employee"
```

The output will list the database partitions used by the employee table. The specific output will depend on the number of database partitions in the database and the number of database partitions in the database partition group that is used by the table space where the employee table was created.

What to do next

After you have verified the installation, you can remove the SAMPLE database to free up disk space. However, it is useful to keep the sample database, if you plan to make use of the sample applications.

Enter the **db2 drop database sample** command to drop the SAMPLE database.

Chapter 25. Taking the first steps after the DB2 Database Partitioning Feature installation

After you install a partitioned database environment, you must perform additional tasks and can optionally perform other tasks.

Setting up a working collective to distribute commands to multiple AIX nodes

In a partitioned database environment on AIX, you can set up a working collective to distribute commands to the set of System p SP workstations that participate in your partitioned database system. Commands can be distributed to the workstations by the **dsh** command.

Before you begin

This can be useful when installing or administering a partitioned database system on AIX, to enable you to quickly execute the same commands on all the computers in your environment with less opportunity for error.

You must know the host name of each computer that you want to include in the working collective.

You must be logged on to the Control workstation as a user with root user authority.

You must have a file that lists the host names for all of the workstations that will participate in your partitioned database system.

Procedure

To set up the working collective to distribute commands to a list of workstations:

1. Create a file called `nodelist.txt` that will list the host names for all of the workstations that will participate in the working collective.

For example, assume that you wanted to create a working collective with two workstations called `workstation1` and `workstation2`. The contents of `nodelist.txt` would be:

```
workstation1
workstation2
```

2. Update the working collective environment variable. To update this list, enter the following command:

```
export DSH_NODE_LIST=path/nodelist.txt
```

where *path* is the location where `nodelist.txt` was created, and `nodelist.txt` is the name of the file that you created that lists the workstations in the working collective.

3. Verify that the names in the working collective are indeed the workstations that you want, by entering the following command:

```
dsh -q
```

You will receive output similar to the following:

```
Working collective file /odelist.txt:
workstation1
workstation2
Fanout: 64
```

Updating the node configuration file (Linux and UNIX)

In a DB2 partitioned database environment, this task provides steps for updating the `db2nodes.cfg` file to include entries for participating computers.

Before you begin

- The DB2 database product must be installed on all participating computers.
- A DB2 instance must exist on the primary computer.
- You must be a user with SYSADM authority.
- Review the configuration examples and file format information provided in the *Format of the DB2 node configuration file* topic if either of the following conditions apply:
 - You plan to use a high speed switch for communication between database partition servers
 - Your partitioned configuration will have multiple logical partitions

About this task

The node configuration file (`db2nodes.cfg`), located in the instance owner's home directory, contains configuration information that tells the DB2 database system which servers participate in an instance of the partitioned database environment. A `db2nodes.cfg` file exists for each instance in a partitioned database environment.

The `db2nodes.cfg` file must contain one entry for each server participating in the instance. When you create an instance, the `db2nodes.cfg` file is automatically created and an entry for the instance-owning server is added.

For example, when you created the DB2 instance using the DB2 Setup wizard, on the instance-owning server `ServerA`, the `db2nodes.cfg` file is updated as follows:

```
0      ServerA      0
```

Restrictions

The hostnames used in the steps of the Procedure section must be fully qualified hostnames.

Procedure

To update the `db2nodes.cfg` file:

1. Log on as the instance owner. For example, `db2inst1` is the instance owner in these steps.
2. Ensure that the DB2 instance is stopped by entering:

```
INSTHOME/sql1lib/adm/db2stop
```

where *INSTHOME* is the home directory of the instance owner (the `db2nodes.cfg` file is locked when the instance is running and can only be edited when the instance is stopped).

For example, if your instance home directory is `/db2home/db2inst1`, enter the following command:

```
/db2home/db2inst1/sqllib/adm/db2stop
```

3. Add an entry to the `.rhosts` file for each DB2 instance. Update the file by adding the following:

```
hostname db2instance
```

where *hostname* is the TCP/IP host name of the database server and *db2instance* is the name of the instance you use to access the database server.

4. Add an entry to the `db2nodes.cfg` file of each participating server. When you first view the `db2nodes.cfg` file, it should contain an entry similar to the following:

```
0      ServerA      0
```

This entry includes the database partition server number (node number), the TCP/IP host name of the server where the database partition server resides, and a logical port number for the database partition server.

For example, if you are installing a partitioned configuration with four computers and a database partition server on each computer, the updated `db2nodes.cfg` should appear similar to the following:

```
0      ServerA      0
1      ServerB      0
2      ServerC      0
3      ServerD      0
```

5. When you have finished updating the `db2nodes.cfg` file, enter the `INSTHOME/sqllib/adm/db2start` command, where *INSTHOME* is the home directory of the instance owner. For example, if your instance home directory is `/db2home/db2inst1`, enter the following command:

```
/db2home/db2inst1/sqllib/adm/db2start
```

6. Log out.

Format of the DB2 node configuration file

The `db2nodes.cfg` file is used to define the database partition servers that participate in a DB2 instance. The `db2nodes.cfg` file is also used to specify the IP address or host name of a high-speed interconnect, if you want to use a high-speed interconnect for database partition server communication.

The format of the `db2nodes.cfg` file on Linux and UNIX operating systems is as follows:

```
dbpartitionnum hostname logicalport netname resourcesetname
```

dbpartitionnum, *hostname*, *logicalport*, *netname*, and *resourcesetname* are defined in the following section.

The format of the `db2nodes.cfg` file on Windows operating systems is as follows:

```
dbpartitionnum hostname computername logicalport netname resourcesetname
```

On Windows operating systems, these entries to the `db2nodes.cfg` are added by the **db2ncrt** or **START DBM ADD DBPARTITIONNUM** commands. The entries can also be modified by the **db2nchg** command. You should not add these lines directly or edit this file.

dbpartitionnum

A unique number, between 0 and 999, that identifies a database partition server in a partitioned database system.

To scale your partitioned database system, you add an entry for each database partition server to the `db2nodes.cfg` file. The *dbpartitionnum* value

that you select for additional database partition servers must be in ascending order; however, gaps can exist in this sequence. You can choose to put a gap between the *dbpartitionnum* values if you plan to add logical partition servers and want to keep the nodes logically grouped in this file.

This entry is required.

hostname

The TCP/IP host name of the database partition server for use by the FCM. This entry is required. Canonical hostname is *strongly* recommended.

If host names are supplied in the *db2nodes.cfg* file, instead of IP addresses, the database manager will dynamically try to resolve the host names. Resolution can be either local or through lookup at registered Domain Name Servers (DNS), as determined by the OS settings on the machine.

Starting with DB2 Version 9.1, both TCP/IPv4 and TCP/IPv6 protocols are supported. The method to resolve host names has changed.

While the method used in pre-Version 9.1 releases resolves the string as defined in the *db2nodes.cfg* file, the method in Version 9.1 or later tries to resolve the Fully Qualified Domain Names (FQDN) when short names are defined in the *db2nodes.cfg* file. Specifying short configured for fully qualified host names, this may lead to unnecessary delays in processes that resolve host names.

To avoid any delays in DB2 commands that require host name resolution, use any of the following work arounds:

1. If short names are specified in the *db2nodes.cfg* files and the operating system host name file, specify the short name and the fully qualified domain name for host name in the operating system host files.
2. To use only IPv4 addresses when you know that the DB2 server listens on an IPv4 port, issue the following command:

```
db2 catalog tcpip4
node db2tcp2 remote 192.0.32.67
server db2inst1 with "Look up IPv4 address from 192.0.32.67"
```

3. To use only IPv6 addresses when you know that the DB2 server listens on an IPv6 port, issue the following command:

```
db2 catalog tcpip6
node db2tcp3 1080:0:0:0:8:800:200C:417A
server 50000
with "Look up IPv6 address from 1080:0:0:0:8:800:200C:417A"
```

logicalport

Specifies the logical port number for the database partition server. This field is used to specify a particular database partition server on a workstation that is running logical database partition servers.

DB2 reserves a port range (for example, 60000 - 60003) in the */etc/services* file for inter-partition communications at the time of installation. This *logicalport* field in *db2nodes.cfg* specifies which port in that range you want to assign to a particular logical partition server.

If there is no entry for this field, the default is 0. However, if you add an entry for the *netname* field, you must enter a number for the *logicalport* field.

If you are using logical database partitions, the *logicalport* value you specify *must* start at 0 and continue in ascending order (for example, 0,1,2).

Furthermore, if you specify a *logicalport* entry for one database partition server, you must specify a *logicalport* for each database partition server listed in your `db2nodes.cfg` file.

This field is optional only if you are *not* using logical database partitions or a high speed interconnect.

netname

Specifies the host name or the IP address of the high speed interconnect for FCM communication.

If an entry is specified for this field, all communication between database partition servers (except for communications as a result of the **db2start**, **db2stop**, and **db2_all** commands) is handled through the high speed interconnect.

This parameter is required only if you are using a high speed interconnect for database partition communications.

resourcesetname

The *resourcesetname* defines the operating system resource that the node should be started in. The *resourcesetname* is for process affinity support, used for Multiple Logical Nodes (MLNs). This support is provided with a string type field formerly known as quadname.

This parameter is only supported on AIX, HP-UX, and Solaris Operating System.

On AIX, this concept is known as "resource sets" and on Solaris Operating System it is called "projects". Refer to your operating systems documentation for more information about resource management.

On HP-UX, the *resourcesetname* parameter is the name of a PRM group. Refer to "HP-UX Process Resource Manager. User Guide. (B8733-90007)" documentation from HP for more information.

On Windows operating systems, process affinity for a logical node can be defined through the **DB2PROCESSORS** registry variable.

On Linux operating systems, the *resourcesetname* column defines a number that corresponds to a Non-Uniform Memory Access (NUMA) node on the system. The system utility **numactl** must be available as well as a 2.6 Kernel with NUMA policy support.

The *netname* parameter must be specified if the *resourcesetname* parameter is used.

Example configurations

Use the following example configurations to determine the appropriate configuration for your environment.

One computer, four database partitions servers

If you are not using a clustered environment and want to have four database partition servers on one physical workstation called ServerA, update the `db2nodes.cfg` file as follows:

0	ServerA	0
1	ServerA	1
2	ServerA	2
3	ServerA	3

Two computers, one database partition server per computer

If you want your partitioned database system to contain two physical workstations, called ServerA and ServerB, update the `db2nodes.cfg` file as follows:

0	ServerA	0
1	ServerB	0

Two computers, three database partition server on one computer

If you want your partitioned database system to contain two physical workstations, called ServerA and ServerB, and ServerA is running 3 database partition servers, update the `db2nodes.cfg` file as follows:

4	ServerA	0
6	ServerA	1
8	ServerA	2
9	ServerB	0

Two computers, three database partition servers with high speed switches

If you want your partitioned database system to contain two computers, called ServerA and ServerB (with ServerB running two database partition servers), and use a high speed interconnect called `switch1` and `switch2`, update the `db2nodes.cfg` file as follows:

0	ServerA	0	switch1
1	ServerB	0	switch2
2	ServerB	1	switch2

Examples using `resourcesetname`

These restrictions apply to the following examples:

- This example shows the usage of `resourcesetname` when there is no high speed interconnect in the configuration.
- The `netname` is the fourth column and a `hostname` also can be specified on that column where there is no switch name and you want to use `resourcesetname`. The fifth parameter is `resourcesetname` if it is defined. The resource group specification can only show as the fifth column in the `db2nodes.cfg` file. This means that for you to specify a resource group, you must also enter a fourth column. The fourth column is intended for a high speed switch.
- If you do not have a high speed switch or you do not want to use it, you must then enter the `hostname` (same as the second column). In other words, the DB2 database management system does not support column gaps (or interchanging them) in the `db2nodes.cfg` files. This restriction already applies to the first three columns, and now it applies to all five columns.

AIX example

Here is an example of how to set up the resource set for AIX operating systems.

In this example, there is one physical node with 32 processors and 8 logical database partitions (MLNs). This example shows how to provide process affinity to each MLN.

1. Define resource sets in `/etc/rset`:

```
DB2/MLN1:
owner      = db2inst1
group      = system
perm       = rwr-r-
resources  = sys/cpu.00000,sys/cpu.00001,sys/cpu.00002,sys/cpu.00003
```

```
DB2/MLN2:
```

```

owner      = db2inst1
group      = system
perm       = rwr-r-
resources  = sys/cpu.00004,sys/cpu.00005,sys/cpu.00006,sys/cpu.00007

```

```

DB2/MLN3:
owner      = db2inst1
group      = system
perm       = rwr-r-
resources  = sys/cpu.00008,sys/cpu.00009,sys/cpu.00010,sys/cpu.00011

```

```

DB2/MLN4:
owner      = db2inst1
group      = system
perm       = rwr-r-
resources  = sys/cpu.00012,sys/cpu.00013,sys/cpu.00014,sys/cpu.00015

```

```

DB2/MLN5:
owner      = db2inst1
group      = system
perm       = rwr-r-
resources  = sys/cpu.00016,sys/cpu.00017,sys/cpu.00018,sys/cpu.00019

```

```

DB2/MLN6:
owner      = db2inst1
group      = system
perm       = rwr-r-
resources  = sys/cpu.00020,sys/cpu.00021,sys/cpu.00022,sys/cpu.00023

```

```

DB2/MLN7:
owner      = db2inst1
group      = system
perm       = rwr-r-
resources  = sys/cpu.00024,sys/cpu.00025,sys/cpu.00026,sys/cpu.00027

```

```

DB2/MLN8:
owner      = db2inst1
group      = system
perm       = rwr-r-
resources  = sys/cpu.00028,sys/cpu.00029,sys/cpu.00030,sys/cpu.00031

```

2. Enable memory affinity by typing the following command:

```
vmo -p -o memory_affinity=1
```

3. Give instance permissions to use resource sets:

```

chuser capabilities=
CAP_BYPASS_RAC_VMM,CAP_PROPAGATE,CAP_NUMA_ATTACH db2inst1

```

4. Add the resource set name as the fifth column in db2nodes.cfg:

```

1 regatta 0 regatta DB2/MLN1
2 regatta 1 regatta DB2/MLN2
3 regatta 2 regatta DB2/MLN3
4 regatta 3 regatta DB2/MLN4
5 regatta 4 regatta DB2/MLN5
6 regatta 5 regatta DB2/MLN6
7 regatta 6 regatta DB2/MLN7
8 regatta 7 regatta DB2/MLN8

```

HP-UX example

This example shows how to use PRM groups for CPU shares on a machine with 4 CPUs and 4 MLNs and 24% of CPU share per MLN, leaving 4% for other applications. The DB2 instance name is db2inst1.

1. Edit GROUP section of /etc/prmconf:

```
OTHERS:1:4::
db2prm1:50:24::
db2prm2:51:24::
db2prm3:52:24::
db2prm4:53:24::
```

2. Add instance owner entry to `/etc/prmconf`:

```
db2inst1:::OTHERS,db2prm1,db2prm2,db2prm3,db2prm4
```
3. Initialize groups and enable CPU manager by entering the following command:

```
prmconfig -i
prmconfig -e CPU
```
4. Add PRM group names as a fifth column to `db2nodes.cfg`:

```
1 voyager 0 voyager db2prm1
2 voyager 1 voyager db2prm2
3 voyager 2 voyager db2prm3
4 voyager 3 voyager db2prm4
```

PRM configuration (steps 1-3) may be done using interactive GUI tool **xprm**.

Linux example

On Linux operating systems, the *resourcesetname* column defines a number that corresponds to a Non-Uniform Memory Access (NUMA) node on the system. The **numactl** system utility must be available in addition to a 2.6 kernel with NUMA policy support. Refer to the man page for **numactl** for more information about NUMA support on Linux operating systems.

This example shows how to set up a four node NUMA computer with each logical node associated with a NUMA node.

1. Ensure that NUMA capabilities exist on your system.
2. Issue the following command:

```
$ numactl --hardware
```

Output similar to the following displays:

```
available: 4 nodes (0-3)
node 0 size: 1901 MB
node 0 free: 1457 MB
node 1 size: 1910 MB
node 1 free: 1841 MB
node 2 size: 1910 MB
node 2 free: 1851 MB
node 3 size: 1905 MB
node 3 free: 1796 MB
```

3. In this example, there are four NUMA nodes on the system. Edit the `db2nodes.cfg` file as follows to associate each MLN with a NUMA node on the system:

```
0 hostname 0 hostname 0
1 hostname 1 hostname 1
2 hostname 2 hostname 2
3 hostname 3 hostname 3
```

Solaris example

Here is an example of how to set up the project for Solaris Version 9.

In this example, there is 1 physical node with 8 processors: one CPU will be used for the default project, three (3) CPUs will be used by the Application Server, and four (4) CPUs for DB2. The instance name is `db2inst1`.

1. Create a resource pool configuration file using an editor. For this example, the file will be called `pool.db2`. Here's the content:

```
create system hostname
create pset pset_default (uint pset.min = 1)
create pset db0_pset (uint pset.min = 1; uint pset.max = 1)
create pset db1_pset (uint pset.min = 1; uint pset.max = 1)
create pset db2_pset (uint pset.min = 1; uint pset.max = 1)
create pset db3_pset (uint pset.min = 1; uint pset.max = 1)
create pset appsrv_pset (uint pset.min = 3; uint pset.max = 3)
create pool pool_default (string pool.scheduler="TS";
    boolean pool.default = true)
create pool db0_pool (string pool.scheduler="TS")
create pool db1_pool (string pool.scheduler="TS")
create pool db2_pool (string pool.scheduler="TS")
create pool db3_pool (string pool.scheduler="TS")
create pool appsrv_pool (string pool.scheduler="TS")
associate pool pool_default (pset pset_default)
associate pool db0_pool (pset db0_pset)
associate pool db1_pool (pset db1_pset)
associate pool db2_pool (pset db2_pset)
associate pool db3_pool (pset db3_pset)
associate pool appsrv_pool (pset appsrv_pset)
```

2. Edit the `/etc/project` file to add the DB2 projects and appsrv project as follows:

```
system:0:::
user:root:1:::
noproject:2:::
default:3:::
group:staff:10:::
appsrv:4000:App Serv project:root::project.pool=appsrv_pool
db2proj0:5000:DB2 Node 0 project:db2inst1,root::project.pool=db0_pool
db2proj1:5001:DB2 Node 1 project:db2inst1,root::project.pool=db1_pool
db2proj2:5002:DB2 Node 2 project:db2inst1,root::project.pool=db2_pool
db2proj3:5003:DB2 Node 3 project:db2inst1,root::project.pool=db3_pool
```

3. Create the resource pool: `# poolcfg -f pool.db2`.
4. Activate the resource pool: `# pooladm -c`
5. Add the project name as the fifth column in `db2nodes.cfg` file:

```
0 hostname 0 hostname db2proj0
1 hostname 1 hostname db2proj1
2 hostname 2 hostname db2proj2
3 hostname 3 hostname db2proj3
```

Enabling communications between database partition servers (Linux and UNIX)

This task describes how to enable communication between the database partition servers that participate in your partitioned database system.

Communication between database partition servers is handled by the Fast Communications Manager (FCM). To enable FCM, a port or port range must be reserved in the `/etc/services` file on each computer in your partitioned database system.

Before you begin

You must have a user ID with root user authority.

You must perform this task on all computers that participate in the instance.

About this task

The number of ports to reserve for FCM is equal to the maximum number of database partitions hosted, or potentially hosted, by any computer in the instance.

In the following example, the `db2nodes.cfg` file contains these entries:

```
0 server1 0
1 server1 1
2 server2 0
3 server2 1
4 server2 2
5 server3 0
6 server3 1
7 server3 2
8 server3 3
```

Assume that the FCM ports are numbered starting at 60000. In this situation:

- server1 uses two ports (60000, 60001) for its two database partitions
- server2 uses three ports (60000, 60001, 60002) for its three database partitions
- server3 uses four ports (60000, 60001, 60002, 60003) for its four database partitions

All computers must reserve 60000, 60001, 60002, and 60003, since this is the largest port range required by any computer in the instance.

If you use a high availability solution such as Tivoli System Automation or IBM PowerHA® SystemMirror for AIX to fail over database partitions from one computer to another, you must account for potential port requirements. For example, if a computer normally hosts four database partitions, but another computer's two database partitions could potentially fail over to it, six ports must be planned for that computer.

When you create an instance, a port range is reserved on the primary computer. The primary computer is also known as the instance-owning computer. However, if the port range originally added to the `/etc/services` file is not sufficient for your needs, you will need to extend the range of reserved ports by manually adding additional entries.

Procedure

To enable communications between servers in a partitioned database environment using `/etc/services`:

1. Log on to the primary computer (instance owning computer) as a user with root authority.
2. Create an instance.
3. View the default port range that has been reserved in the `/etc/services` file. In addition to the base configuration, the FCM ports should appear similar to the following:

```
db2c_db2inst1      50000/tcp
#Add_FCM port information
DB2_db2inst1       60000/tcp
DB2_db2inst1_1     60001/tcp
DB2_db2inst1_2     60002/tcp
DB2_db2inst1_END   60003/tcp
```

By default, the first port (50000) is reserved for connection requests, and the first available four ports above 60000 are reserved for FCM communication.

One port is for the instance-owning database partition server and three ports are for logical database partition servers that you might choose to add to the computer after installation is complete.

The port range must include a start and an END entry. Intermediate entries are optional. Explicitly including intermediate values can be useful for preventing other applications from using these ports, but these entries are not verified by the database manager.

DB2 port entries use the following format:

```
DB2_instance_name_suffix port_number/tcp # comment
```

where:

- *instance_name* is the name of the partitioned instance.
 - *suffix* is not used for the first FCM port. Intermediate entries are those between the lowest and highest port. If you include the intermediate entries between the first and ending FCM port, the *suffix* consists of an integer that you increment by one for each additional port. For example, the second port is numbered 1, and third is numbered 2, and so on to ensure uniqueness. The word END must be used as the *suffix* for the last entry.
 - *port_number* is the port number that you reserve for database partition server communications.
 - *comment* is an optional comment describing an entry.
4. Ensure that there are sufficient ports reserved for FCM communication. If the range of reserved ports is insufficient, add new entries to the file.
 5. Log on as a root user to each computer participating in the instance and add identical entries to the /etc/services file.

Enabling the execution of remote commands (Linux and UNIX)

You must update your .rhosts file to execute remote commands using **rsh**.

Before you begin

In a partitioned database system, each database partition server must have the authority to perform remote commands on all the other database partition servers participating in an instance. This authority can be granted by updating the .rhosts file in the home directory for the instance. Because the home directory for the instance is on the shared DB2 home file system, only one .rhosts file is required.

- You must have root user authority.
- You must know the host name of each participating computer.
- You must know the instance owner's user name.

About this task

This topic describes how to enable execution of remote commands using **rsh**.

You can also use **ssh** to enable execution of remote commands. To use **ssh** without being prompted for passwords or pass phrases, refer to:

- Chapter 22, "Setting up a partitioned database environment," on page 319
- <http://www.ibm.com/developerworks/data/library/techarticle/dm-0506finnie/index.html>

Procedure

To update your `.rhosts` file to execute remote commands using `rsh`:

1. Log onto the primary computer as a user with root user authority.
2. Create a `.rhosts` file in the instance home directory. For example, if your instance home directory is `/db2home/db2inst1`, you can use a text editor to create the `.rhosts` file by entering the following command:

```
vi /db2home/db2inst1/.rhosts
```

3. Add entries to the `.rhosts` file for each computer including the primary computer. The `.rhosts` file has the following format:

```
hostname instance_owner_user_name
```

Some systems might require a long host name to be specified, for example: `ServerA.yourdomain.com`. Before you add host name entries to the `.rhosts` file, make sure the host names in the `/etc/hosts` and the `/etc/resolv.conf` files can be resolved.

The `INSTHOME/.rhosts` file should contain entries similar to the following:

```
ServerA.yourdomain.com db2inst1
ServerB.yourdomain.com db2inst1
ServerC.yourdomain.com db2inst1
ServerD.yourdomain.com db2inst1
```

Rather than specifying each host name individually, you can specify the following entry in the `.rhosts` file, but this action may pose a security risk and should only be done in a test environment.

```
+ db2inst1
```

If you have specified a high speed switch (netname) in the `db2nodes.cfg` file, you should also add netname entries for each computer to the `.rhosts` file. The netname values are specified in the fourth column of the `db2nodes.cfg` file. A `.rhosts` file with high speed switch (netname) entries may look similar to the following:

```
ServerA.yourdomain.com db2inst1
ServerB.yourdomain.com db2inst1
ServerC.yourdomain.com db2inst1
ServerD.yourdomain.com db2inst1
Switch1.yourdomain.com db2inst1
Switch2.yourdomain.com db2inst1
Switch3.yourdomain.com db2inst1
Switch4.yourdomain.com db2inst1
```

What to do next

An alternative to using a `.rhosts` file is to use `/etc/hosts.equiv` file. The `/etc/hosts.equiv` file would contain the exact same entries as the `.rhosts` file, but must be created on each computer.

For more information about the `.rhosts` file or the `/etc/hosts.equiv` file, see your operating system documentation.

Chapter 26. Response files

Using the response file that you created by using the DB2 Setup wizard, you can install database partition servers on a different operating system.

Installing database partition servers on participating computers using a response file (Windows)

In this task you will use the response file you created using the DB2 Setup wizard to install database partition servers on participating computers.

Before you begin

- You have installed a DB2 copy on the primary computer using the DB2 Setup wizard.
- You have created a response file for installing on participating computers and copied it onto the participating computer.
- You must have administrative authority on participating computers.

Procedure

To install additional database partition servers using a response file:

1. Log to the computer that will participate in the partitioned database environment with the local Administrator account that you have defined for the DB2 installation.
2. Change to the directory containing the DB2 database product DVD. For example:

```
cd c:\db2dvd
```

where db2dvd represents the name of the directory containing the DB2 database product DVD.

3. From a command prompt, enter the **setup** command as follows:
`setup -u responsefile_directory\response_file_name`

In the following example, the response file, Addpart.file can be found in the c:\responsefile directory. The command for this example, would be:

```
setup -u c:\responsefile\Addpart.file
```

4. Check the messages in the log file when the installation finishes. You can find the log file in the My Documents\DB2LOG\ directory. You should see output similar to the following at the end of the log file:

```
=== Logging stopped: 5/9/2007 10:41:32 ===  
MSI (c) (C0:A8) [10:41:32:984]: Product: DB2  
Enterprise Server Edition - DB2COPY1 -- Installation  
operation completed successfully.
```

5. When you install the instance-owning database partition server on the primary computer, the DB2 database product reserves a port range according to the specified number of logical database partition servers participating in partitioned database environment. The default range is four ports. For each server that participates in the partitioned database environment, you must manually configure the /etc/services file for the FCM ports. The range of the FCM ports depends on how many logical partitions you want to use on the

participating computer. A minimum of two entries are required, *DB2_instance* and *DB2_instance_END*. Other requirements for the FCM ports specified on participating computers are:

- The starting port number must match the starting port number of the primary computer.
- Subsequent ports must be sequentially numbered.
- Specified port numbers must be free.

Results

You must log onto each participating computer and repeat these steps.

What to do next

If you want your DB2 database product to have access to DB2 documentation either on your local computer or on another computer on your network, then you must install the *DB2 Information Center*. The *DB2 Information Center* contains documentation for the DB2 database system and DB2 related products.

Installing database partition servers on participating computers using a response file (Linux and UNIX)

In this task you will use the response file you created using the DB2 Setup wizard to install database partition servers on participating computers.

Before you begin

- You have installed DB2 database product on the primary computer using the DB2 Setup wizard and have created a response file for installing on participating computers.
- You must have root user authority on participating computers.

Procedure

To install additional database partition servers using a response file:

1. As root, log on to a computer that will participate in the partitioned database environment.
2. Change to the directory where you copied the contents of the DB2 database product DVD. For example:

```
cd /db2home/db2dvd
```
3. Enter the **db2setup** command as follows:

```
./db2setup -r /responsefile_directory/response_file_name
```

In this example, the response file, *AddPartitionResponse.file*, was saved to the */db2home* directory. The command for this situation would be:

```
./db2setup -r /db2home/AddPartitionResponse.file
```

4. Check the messages in the log file when the installation finishes.

Results

You must log onto each participating computer and perform a response file installation.

What to do next

If you want your DB2 database product to have access to DB2 database documentation either on your local computer or on another computer on your network, then you must install the *DB2 Information Center*. The *DB2 Information Center* contains documentation for the DB2 database system and DB2 database related products.

Part 5. Response file considerations

There are some considerations that you need to take into account when using response files to install DB2 products for Linux, UNIX, and Windows.

You should understand the following considerations before proceeding with a response file installation:

- Although response files created in Version 9 and response files created in Version 10 have similar formats, there are version limitations as to where response files can be used. For example, any response file generated in DB2 Version 10 can only be used to install a DB2 Version 10 product (for example Version 10.1); the response file cannot be used to install DB2 Version 9. The opposite is also true, where response files generated in DB2 Version 9 cannot be used to install DB2 Version 10. This is primarily caused by mandatory keywords that are new in Version 10.
- On Linux or UNIX platforms, a response file created for a root installation might not be usable for a non-root installation. Some response file keywords are valid for root installation only.
- If you are using the DB2 Setup wizard:
 - You can save your settings in a response file during the installation in the **Select the installation action** panel of the DB2 Setup wizard.
 - You are creating a response file based on just the installation you are performing. This method is recommended if you have either a fairly simple configuration or if you want to create a response file that you plan to later customize.
 - A response file is only generated if you allow the installation process to complete, and it completes successfully. If you cancel the installation, or if the installation fails, the response file is not created.
 - Response files created with this method cannot be used in the **db2i setup** command **-r** parameter without modification. A response file created using the setup wizard, and used in the **db2i setup** command must be modified to meet the following conditions:
 - must contain the keyword **FILE**
 - must not contain the keywords: **PROD**, **LIC_AGREEMENT**, or **INSTALL_TYPE**.
- You can use a response file to install an identical configuration across every workstation on your network or to install multiple configurations of a DB2 database product. You can then distribute this file to every workstation where you want this product to be installed.
- If you use the response file generator, you are creating the response file based on an existing installation. This method is recommended when you have a more complex configuration, one that you manually configured. If you are using the response file generated by the response file generator, you might need to input user names and passwords.

Chapter 27. Response file installation basics

A DB2 response file installation lets you install DB2 products without any user interaction.

A *response file* is an English-only text file that contains setup and configuration information. A response file specifies configuration and setup parameters and the products and components to install.

This method is useful not only for large-scale deployments of DB2 products, but also for embedding the DB2 installation process transparently inside your customized installation and configuration procedure.

You can create a response file by any of the following methods:

- Modifying the `db2dsf.rsp` sample response files provided in the `db2/platform/samples` directory, where *platform* refers to the appropriate operating system.
- Using the DB2 Setup wizard to save the setup and configuration data according to the input you provided. If you choose the option to create a response file in the DB2 Setup wizard, the response files will be saved by default at this location. By default, response files are saved to: `/tmp`.
- In a DB2 pureScale environment a response file can be saved from the instance setup wizard (the **db2isetup** command.)

A response file installation can also be referred to as a silent installation or an unattended installation.

Chapter 28. Response file uninstall basics

You can use a response file to uninstall one or more DB2 database products, features, or languages. A response file uninstall can also be referred to as a silent uninstall or an unattended uninstall.

On Linux and UNIX operating systems, you can also use a response file to uninstall the *DB2 Information Center*.

A sample uninstall response file, `db2un.rsp`, is provided. When the DB2 database product was installed, as part of the installation the sample uninstall response file is copied to `DB2DIR/install`, where `DB2DIR` specifies the path where the DB2 database product was installed. You can customize the response file. To activate a keyword in the response file, remove the asterisk (*) to the left of the keyword.

Before uninstalling using a response file, the following considerations, and restrictions, must be considered.

When removing a DB2 database product, you must first drop the DB2 instance. For example, if only DB2 ESE is installed, before removing the ESE product you must first drop the instance.

If multiple DB2 database products exist within the same DB2 copy, a response file uninstallation of one DB2 database product does not affect the components shared by other DB2 database products in that copy. For example, DB2COPY1 contains the DB2 database products ESE, WSE and PE. Several components are shared among the three products. In this case, uninstalling ESE removes the components that are not shared. In the following figure, the shaded region represents the components to be removed from DB2COPY1:

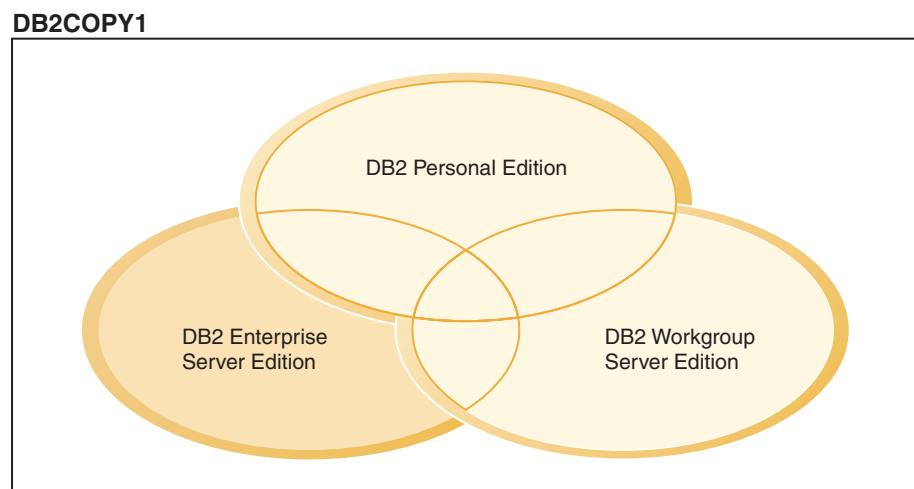


Figure 45. Shared components among different DB2 database products within the same DB2 copy

However, before removing a DB2 database product, consider the instance type. For example, ESE, WSE and PE are installed in the same DB2 copy. You want to remove the ESE product. Before removing the ESE product, you can either:

- move the instance to another DB2 copy of the same version that supports the ESE instance type using the **db2iupdt** command, or,
- remove the existing ESE instance using the **db2idrop** command.

If a DB2 feature is installed by multiple products in the same DB2 copy, a response file uninstallation of the feature removes the feature from all of the products in the DB2 copy.

The following restrictions apply:

- A DB2 database product cannot be removed if that product is required by an add-on product. For example, if both ESE and QP are installed, ESE cannot be uninstalled.
- A DB2 feature cannot be removed if that feature is required by another feature.
- The language English cannot be removed.
- On Windows operating systems:
 - When you are removing a DB2 database product, the uninstall response file cannot also contain the keywords to remove a feature or a language. The keyword **REMOVE_PROD** cannot be combined with keywords **REMOVE_COMP** or **REMOVE_LANG**.
 - If you have a clustered environment, before uninstalling your DB2 database product you must run the **db2mscs** command, with the **-u** option, from the same server that originally ran the **db2mscs** command to create the failover infrastructure. For details, see the **db2mscs** command.

After a response file uninstall, check the log file. This file captures all DB2 uninstallation information, including errors.

- On Linux and UNIX operating systems, the log file is located in `/tmp/db2_deinstall-nnnnn.log` where *nnnnn* are generated numbers. The log file name displays on the screen after running the **db2_deinstall** command. You can also verify the product, feature, or language was removed using the **db2ls** command.
- On Windows operating systems, the log file is located in `My Documents\DB2LOG\db2un-TimeStamp.log`. You can verify the product, feature, or language was removed by opening the Add/Remove Programs dialog in the control panel.

Chapter 29. Response file considerations for a DB2 pureScale environment

In addition to general response file considerations, a DB2 pureScale environment has additional considerations.

A response file installation might also be referred to as a silent installation or an unattended installation. The response file can be used to:

- Install the DB2 pureScale Feature and create the DB2 pureScale instance in the same run across all hosts, or
- Install the DB2 pureScale Feature across all hosts, then run the **db2icrt**, **db2iupdt**, or **db2isetup** commands to manually create and extend a DB2 pureScale instance.

A response file can be created in either of the following ways:

- Using the DB2 Setup wizard,
- using the sample response file, or
- using the **db2isetup** command.

The DB2 pureScale Feature sample response file, `db2dsf.rsp`, is located in `/opt/IBM/db2/V10.1/install/db2/platform/samples` directory, where *platform* refers to the appropriate operating system.

You must consider the following before proceeding with a response file installation in a DB2 pureScale environment:

- A clean system is recommended. In a clean system, neither of the following products is already installed:
 - IBM General Parallel File System (GPFS), nor
 - IBM Tivoli System Automation for Multiplatforms
- The response file includes information about all the participating hosts. When using a response file, the **db2setup** or **db2isetup** command can be run from any of the hosts. For additional usage notes, see the specific command.
- Any keyword with incorrect syntax results in the entire response file being invalid.
- Keywords available for a DB2 pureScale Feature installation. Some keywords are mandatory, some are optional depending on the state of your system. For keyword details, see the Response file keywords topic.
 - **HOST**
 - *host_identifier*.**HOSTNAME**
 - *host_identifier*.**CLUSTER_INTERCONNECT_NETNAME**
 - **DB2_CLUSTER_SERVICES_TIEBREAKER_DEVICE_PATH**
 - **INSTANCE_SHARED_DEVICE_PATH**
 - **INSTANCE_SHARED_DIR**
 - Instance creation settings:
 - *instance_identifier*.**PREFERRED_PRIMARY_CF**
 - *instance_identifier*.**PREFERRED_SECONDARY_CF**
 - *instance_identifier*.**MEMBER**
 - *instance_identifier*.**DB2CF_PORT_NUMBER**

- *instance_identifier.DB2CF_MGMT_PORT_NUMBER*
- The following keyword interdependencies must be adhered to when using the keywords:
 - The recommended DB2 pureScale environment setup for the DB2 pureScale Feature is four hosts: two cluster caching facilities and two members. For example:


```

* Product Installation
LIC_AGREEMENT = ACCEPT
PROD = ENTERPRISE_SERVER_EDITION_DSF
FILE = /opt/IBM/db2/V9.8
* -----
* Instance properties
* -----
INSTANCE = inst1
inst1.PREFERRED_PRIMARY_CF = host1
inst1.PREFERRED_SECONDARY_CF = host2
inst1.MEMBER = host3
inst1.MEMBER = host4
inst1.TYPE = dsf
* Instance-owning user
inst1.NAME = db2sdin1
inst1.GROUP_NAME = db2iadm1
inst1.HOME_DIRECTORY = /home/db2sdin1
inst1.PASSWORD = 937436363038076535953807826033486554358563633346043342
ENCRYPTED = inst1.PASSWORD
inst1.START_DURING_INSTALL = YES
inst1.SVCENAME = db2c_db2sdin1
inst1.PORT_NUMBER = 50000
inst1.FCM_PORT_NUMBER = 60000
inst1.MAX_LOGICAL_NODES = 1
inst1.DB2CF_MGMT_PORT_NUMBER = 56000
inst1.DB2CF_PORT_NUMBER = 56001
* Fenced user
inst1.FENCED_USERNAME = db2sdfe1
inst1.FENCED_GROUP_NAME = db2fadm1
inst1.FENCED_HOME_DIRECTORY = /home/db2sdfe1
inst1.FENCED_PASSWORD = 786359555684387445670627543570638933213376259440013362
ENCRYPTED = inst1.FENCED_PASSWORD
*-----
* Installed Languages
*-----
LANG = EN
*-----
* Host Information
*-----
HOST = host1
host1.HOSTNAME = coralpibinstall05
host1.CLUSTER_INTERCONNECT_NETNAME = coralpibinstall05-ib0
HOST = host2
host2.HOSTNAME = coralpibinstall06
host2.CLUSTER_INTERCONNECT_NETNAME = coralpibinstall06-ib0
HOST = host3
host3.HOSTNAME = coralpibinstall07
host3.CLUSTER_INTERCONNECT_NETNAME = coralpibinstall07-ib0
HOST = host4
host4.HOSTNAME = coralpibinstall08
host4.CLUSTER_INTERCONNECT_NETNAME = coralpibinstall08-ib0

* -----
* Shared file system settings
* -----
INSTANCE_SHARED_DEVICE_PATH = /dev/hdisk10
* -----
          
```



```
* Tiebreaker settings
* -----
DB2_CLUSTER_SERVICES_TIEBREAKER_DEVICE_PATH = /dev/hdisk11
```

where:

- *instance_identifier*.**PREFERRED_PRIMARY_CF** - identifies the host designated as the preferred primary cluster caching facility for the instance. One host must be specified as the preferred primary cluster caching facility.
 - *instance_identifier*.**PREFERRED_SECONDARY_CF** - identifies the host designated as the secondary cluster caching facility for the instance.
 - *instance_identifier*.**MEMBER** - identifies the host designated as a member for the instance.
- File system keywords. When the DB2 pureScale instance is created, one of the cluster file system related keywords must be specified. The keyword chosen indicates whether an existing User-Managed GPFS file system, or a created GPFS file system is used:
- **INSTANCE_SHARED_DEVICE_PATH** - indicates a DB2 Managed clustered file system method is used. This is the recommended keyword. This keyword identifies the device path used to create a file system for shared files. There can be no file system already on the disk.
 - **INSTANCE_SHARED_DIR** - indicates a User-Managed file system method is used. This keyword identifies the path to an existing GPFS file system for storing shared files.

Only one of the file system keywords should be specified, not both. If both keywords are specified, **INSTANCE_SHARED_DIR** is ignored. However, if the global registry record **DEFAULT_INSTPROF** contains a value, this means there is an existing file system. (This can happen if a DB2 pureScale instance was created and then dropped. For details see the **db2idrop** command.) If the global registry record **DEFAULT_INSTPROF** contains a value, neither keyword needs to be specified.

- The **DB2_CLUSTER_SERVICES_TIEBREAKER_DEVICE_PATH** keyword is the RSCT tiebreaker setting. It identifies the device path to setup a cluster services tiebreaker. The tiebreaker is required for cluster recovery purposes when a DB2 pureScale instance is split in half because of a communication failure. The tiebreaker determines which half of the DB2 pureScale instance remains online. The path specified for this keyword should be dedicated, and not used for anything else. It should be at least 25 MBs in size, and accessible by all hosts. The path specified for this keyword must be different from the **INSTANCE_SHARED_DEVICE_PATH** keyword.

Example

The following example installs the DB2 pureScale Feature using response file `db2dsf.rsp`:

```
db2setup -r /mydir/db2dsf.rsp -l /tmp/db2setup.log -t /tmp/db2setup.trc
```

where `mydir` represent the directory where you saved the `db2dsf.rsp` file.

Chapter 30. Creating a response file using the DB2 Setup wizard or the DB2 Instance Setup wizard

You can create a response file using the DB2 Setup wizard or the DB2 Instance Setup wizard. You can save your panel selections from the GUI in a response file without having to actually perform a DB2 database product installation.

Procedure

- To create a response file using the DB2 Setup wizard:
 1. Type the **db2setup** command to launch the DB2 Instance Setup wizard.
 2. On the Select installation, response file creation, or both panel, select the **Save my installation settings in a response file** option. Specify a location where the DB2 Instance Setup wizard will copy the generated response file. Click **Next**.
 3. Proceed through the remaining panels making the appropriate selections.
 4. On the Start copying files and create response file panel, click **Finish** to generate the response file.
- To create a response file using the DB2 Instance Setup wizard:
 1. Type the **db2isetup** command to launch the DB2 Instance Setup wizard.
 2. On the Select installation, response file creation, or both panel, select the **Save my installation settings in a response file** option. Specify a location where the DB2 Instance Setup wizard will copy the generated response file. Click **Next**.
 3. Proceed through the remaining panels making the appropriate selections.
 4. On the Start copying files and create response file panel, click **Finish** to generate the response file.

What to do next

You can now use the generated response file to perform an unattended installation using the same settings.

Chapter 31. Response file installation of DB2 overview (Linux and UNIX)

This task describes how to perform response file installations on Linux or UNIX. You can use the response file to install additional components or products after an initial installation. A response file installation might also be referred to as a silent installation or an unattended installation.

Before you begin

Before you begin the installation, ensure that:

- Your system meets all of the memory, hardware, and software requirements to install your DB2 database product.
- All DB2 processes are stopped. If you are installing a DB2 database product on top of an existing DB2 installation on the computer, you must stop all DB2 applications, the DB2 database manager, and DB2 processes for all DB2 instances and DB2 DAS related to the existing DB2 installation.

Restrictions

Be aware of the following limitations when using the response files method to install DB2 on Linux or UNIX operating systems:

- If you set any instance or global profile registry keywords to BLANK (the word "BLANK"), that keyword is, in effect, deleted from the list of currently set keywords.
- Ensure that you have sufficient disk space before installing. Otherwise, if the installation fails, manual cleanup is required.
- If you are performing multiple installations or are installing DB2 database products from multiple DVDs, it is recommended that you install from a network file system rather than a DVD drive. Installing from a network file system significantly decreases the amount of time it takes to perform the installation.
- If you are planning on installing multiple clients, set up a mounted file system on a code server to improve performance.

Procedure

To perform a response file installation:

1. Mount your DB2 database product DVD or access the file system where the installation image is stored.
2. Create a response file by using the sample response file.
Response files have a file type of `.rsp`. For example, `ese.rsp`.
3. Install DB2 using the response file.

Creating a response file using the sample response file

After creating a response file, you can make changes to the response file to activate or deactivate keywords.

About this task

One of the following scenarios apply:

- You created a response file using the DB2 Setup wizard (based on your installation) and want to make changes to it.
- You want to create a response file based on a sample response file (db2ese.rsp located in db2/*platform*/samples directory, where *platform* refers to the appropriate operating system).

If you are installing directly from the DVD, you must store the updated or customized response file on another file system.

You can specify the name of the instance owner and the hosts to be part of a DB2 pureScale instance in the response file. If this user does not already exist, the DB2 database product installer creates this user on your system and shared disks.

Procedure

To edit a response file:

1. If you are using the sample response file, copy the sample response file to a local file system and edit it using the *Response file keywords* topic as a guide.
2. Customize the sample response file.

To activate an item in the response file, remove the asterisk (*) to the left of the keyword. Then, replace the current setting to the right of the value with the new setting. The possible settings are listed to the right of the equal sign.

Keywords that are unique to installation are only specified in a response file during a response file installation.

Installing a DB2 database product using a response file (Linux and UNIX)

These instructions guide you through the process of installing a DB2 database product on a Linux or UNIX operating system using a response file.

Before you begin

Before you begin the installation:

- Ensure that your system meets installation, memory, and disk requirements.
- Complete the pre-installation tasks.
- Complete the IBM DB2 pureScale Feature installation checklist.
- Ensure that you are logged in as root.
- For root installations, log on with the user ID that has root privileges. For non-root installations, log on with the user ID that is to own the DB2 installation.
- Your system meets all of the memory, hardware, and software requirements to install your DB2 database product.
- You can install a DB2 database product using a response file with either root or non-root authority.
- All DB2 processes associated with the copy of DB2 you are working with are stopped.
- Ensure the response file keyword LIC_AGREEMENT is set to ACCEPT.

Procedure

To install the DB2 pureScale Feature using a response file:

1. Enter the **db2setup** command as follows:

```
./db2setup -r responsefile_directory/response_file
```

where:

- *responsefile_directory* represents the directory where the response file is located; and
 - *response_file* represents the name of the response file.
2. Check the messages in the log file when the installation finishes. The installation logs are located, by default, in the /tmp directory:
 - db2setup.log (db2setup_username.log for non-root installations, where *username* is the non-root user ID under which the installation was performed)
 - db2setup.err (db2setup_username.err for non-root installations, where *username* is the non-root user ID under which the installation was performed)

You can specify the location of the log file. An installation history log `db2install.history` is located in `DB2DIR/install/logs` where *DB2DIR* specifies the path where the DB2 database product was installed. The default installation directory is:

- For AIX, HP-UX or Solaris `/opt/IBM/db2/V10.1`
- For Linux `/opt/ibm/db2/V10.1`

If more than one DB2 database product is installed in the same location, you will see `db2install.history.xxxx`, where *xxxx* are digits, starting at 0001 and increasing depending on how many DB2 copies you have installed.

Example

The following example installs the DB2 pureScale Feature using response file `db2dsf.rsp`:

```
db2setup -r /mydir/db2dsf.rsp -l /tmp/db2setup.log -t /tmp/db2setup.trc
```

where *mydir* represent the directory where you saved the `db2dsf.rsp` file.

What to do next

If you want your DB2 database product to have access to DB2 documentation either on your local computer or on another computer on your network, then you must install the *DB2 Information Center*. The *DB2 Information Center* contains documentation for the DB2 database and related products.

Review and perform any necessary post-installation tasks.

Installing database partition servers on participating computers using a response file (Linux and UNIX)

In this task you will use the response file you created using the DB2 Setup wizard to install database partition servers on participating computers.

Before you begin

- You have installed DB2 database product on the primary computer using the DB2 Setup wizard and have created a response file for installing on participating computers.
- You must have root user authority on participating computers.

Procedure

To install additional database partition servers using a response file:

1. As root, log on to a computer that will participate in the partitioned database environment.

2. Change to the directory where you copied the contents of the DB2 database product DVD. For example:

```
cd /db2home/db2dvd
```

3. Enter the **db2setup** command as follows:

```
./db2setup -r /responsefile_directory/response_file_name
```

In this example, the response file, `AddPartitionResponse.file`, was saved to the `/db2home` directory. The command for this situation would be:

```
./db2setup -r /db2home/AddPartitionResponse.file
```

4. Check the messages in the log file when the installation finishes.

Results

You must log onto each participating computer and perform a response file installation.

What to do next

If you want your DB2 database product to have access to DB2 database documentation either on your local computer or on another computer on your network, then you must install the *DB2 Information Center*. The *DB2 Information Center* contains documentation for the DB2 database system and DB2 database related products.

Response file error codes (Linux and UNIX)

During installation and uninstallation on Linux and UNIX platforms using response files, you may encounter error code values.

The following tables describe error return codes (primary and secondary) that can be encountered during a response file installation.

Table 47. Primary response file installation error codes

Error code value	Description
0	The action completed successfully.
1	The action returns a warning.
67	A fatal error occurred during the installation.

Table 48. Secondary response file installation error codes

Error code value	Description
3	The path was not found.
5	Access was denied.
10	An environment error occurred.
13	The data is invalid.
16	The DAS creation failed.
17	The instance creation failed.
18	The database creation failed.
19	The fix pack image is at a lower level than the installed product.
20	The fix pack image is at the same level as the installed product. There is nothing to update.
21	There are no DB2 products installed at the selected location. There is nothing to update.
22	To update an already installed DB2 product, use the installFixPack command instead of the db2setup command.
87	One of the parameters was invalid.
66	The installation was cancelled by the user.
74	The configuration data is corrupt. Contact your support personnel.
76	The installation source for this product is not available. Verify that the source exists and that you can access it.
82	Another installation is already in progress. Complete that installation first before proceeding with this installation.
86	There was an error opening the installation log file. Verify that the specified log file location exists and that it is writable.
97	This installation package is not supported on this platform.
102	Another version of this product is already installed. Installation of this version cannot continue.
103	Invalid command line argument.
143	The system does not have enough free space to continue with the installation.

Uninstalling a DB2 database product, feature, or language using a response file (Linux and UNIX)

To silently uninstall DB2 database products, features, or languages in a DB2 copy, use the **db2_deinstall** command with the **-r** option.

About this task

You can also use a response file to uninstall the *DB2 Information Center*.

If you have DB2 database products installed in various installation paths, you must run this command separately from each installation path. Sample response file, **db2un.rsp**, is in *DB2DIR/install*, where *DB2DIR* specifies the path where the DB2 database product was installed.

Procedure

To perform an uninstallation:

1. Customize the `db2un.rsp` response file. To activate an item in the response file, remove the asterisk (*) to the left of the keyword. Then, replace the current setting to the right of the value with the new setting. The possible settings are listed to the right of the equal sign.
2. Run the **`db2_deinstall`** command. For example, **`db2_deinstall -r db2un.rsp`**.
3. Check the messages in the log file when the uninstallation finishes. The log file is located in:
 - For root installations: `/tmp/db2_deinstall.log.process-id`
 - For non-root installations: `/tmp/db2_deinstall_user-id.log`

Chapter 32. Response file installation of DB2 overview (Windows)

On Windows, you can perform a response file installation of a DB2 product on a single machine or on multiple machines. A response file installation might also be referred to as a silent installation or an unattended installation.

Before you begin

Before you begin the installation, ensure that:

- Your system meets all of the memory, hardware, and software requirements to install your DB2 product.
- You have all of the required user accounts to perform the installation.
- Ensure all DB2 processes are stopped.

Procedure

- To perform a response file installation of a DB2 product on a single machine:
 1. Create and customize a response file by one of the following methods:
 - Modifying a sample response file. Sample response files are located in (db2\Windows\samples).
 - Using the DB2 Setup wizard to generate a response file.
 - Using the response file generator.
 2. Run the **setup -u** command specifying your customized response file. For example, a response file created during an installation:

```
setup -u my.rsp
```
- To perform a response file installation of a DB2 product on multiple machines:
 1. Set up shared access to a directory.
 2. Create a response file using the sample response file.
 3. Install a DB2 product using a response file.

Making the DB2 installation files available for a response file installation (Windows)

Making the DB2 installation file available for an installation is part of the process to perform a response file installation of a DB2 database product on multiple machines.

About this task

To make the DB2 installation files available for a response file installation, you must copy the required files from the product DVD to another drive.

Procedure

To copy the required files from the product DVD to another drive:

1. Insert the appropriate DB2 database product DVD into your DVD drive.
2. Create a directory (for example, c:\db2prods).
3. Copy the DB2 installation files.

Setting up shared access to a directory (Windows)

Setting up shared access to a directory is part of the process to perform a response file installation of a DB2 database product on multiple machines. This allows you to grant your network workstations access to a directory on the installation server.

Procedure

To set up shared access to a directory on the installation server:

1. Copy the DB2 installation files to a directory. For example, c:\db2prods.
2. Open Windows Explorer.
3. Select the directory on the installation server that you want to share.
4. Select **File > Properties** from the menu bar. The Properties window for the directory opens.
5. Select the **Sharing** tab.
6. For Windows 2003, select the **Shared this folder** button.
7. For Windows 2008 or Windows Vista (or higher), click the **Advanced Sharing** button. Enable the **Shared this folder** check box.
8. In the **Share Name** field, type a share name. For example, db2win.
9. To specify *Read access* for everyone:
 - a. Click the **Permissions** push button. The Permissions for share_name window opens.
 - b. Ensure that the **Everyone** option is selected in the **Group or user names** box.
 - c. In the **Permissions for Everyone** box, ensure that read permission is enabled. Click the check box in the **Allow** column for **Read**.
 - d. Click **OK**. You are returned to the Properties window of the directory for which you want to set up shared access.
 - e. Click **OK**.

Editing a response file (Windows)

After creating a response file, you can make changes to the response file to activate or deactivate keywords.

About this task

One of the following scenarios apply:

- You want to create a response file based on a sample response file (located in db2\Windows\samples).
- You created a response file using the DB2 Setup wizard (based on your installation) and want to make changes to it.
- You have already set up and configured your DB2 database product and you want to distribute this exact configuration across your network using the response file generated by the response file generator. If you are using the response file generated by the response file generator, you can input user names and passwords.

Restrictions

Install DB2 database products only on a drive which is local to the target workstation. Installing on a non-local drive can cause performance and availability problems.

Procedure

To edit a response file:

1. If you are using the sample response file, make a copy of it and open it in a text editor. If you are using the response file created by the DB2 Setup wizard, open it in a text editor.
2. Customize the response file.
To activate an item in the response file, remove the asterisk (*) to the left of the keyword. Then, replace the current setting to the right of the value with the new setting. The possible settings are listed to the right of the equal sign.
Some product response files have mandatory keywords that you must provide values for. The mandatory keywords are documented in the comments of each response file.
3. Save the file on the shared network drive so that it is available to the installation server. If you have made any changes, save the file under a new file name to preserve the original sample response file. If you are installing directly from the product DVD, store the renamed response file on another drive.

Installing a DB2 product using a response file (Windows)

These instructions will guide you through the process of installing a DB2 database product on a Windows operating system using a response file.

Before you begin

Before you begin the installation, ensure that:

- Your system meets all of the memory, hardware, and software requirements to install your DB2 product.
- Ensure the response file keyword LIC_AGREEMENT is set to ACCEPT.

Procedure

To perform an installation from the workstation where the DB2 product will be installed:

1. Connect to the shared directory of the network drive or DVD drive containing the DB2 installation files by entering the following command from the command prompt:

```
net use x: \\computer_name\directory_sharename /USER:domain\username
```

where:

- *x*: represents the shared directory on the local drive.
- *computer_name* represents the computer name of the remote server where the DB2 installation files are located.
- *directory_sharename* represents the share name of the directory on the network drive or DVD drive where the DB2 installation files are located.
- *domain* represents the domain where the account is defined.
- *username* represents a user that has access to this system.

For example, to use the remote db2prods directory, which was shared as DB2_Installation_Images and is located on the remote server Server, as the local x: drive, enter the following command:

```
net use x: \\Server\DB2_Installation_Images
```

Depending on how security is set up across your network, you might have to specify the **/USER** parameter.

2. Enter the **setup** command as follows:

```
setup -u response_file
```

where *response_file* represents the full path and file name of the response file to use.

If you are using a response file that was created using the response file generator, ensure that all the instance profiles are located in the same drive and directory as the response file that you specify.

3. Check the messages in the log file when the installation finishes. For information about errors encountered during installation, review the installation log file located in the My Documents\DB2LOG\ directory. The log file uses the following format: DB2-ProductAbbrev-DateTime.log, for example, DB2-ESE-Tue Apr 04 17_04_45 2007.log.

What to do next

If you want your DB2 product to have access to DB2 documentation either on your local computer or on another computer on your network, you must install the DB2 Information Center.

Installation of DB2 products using Microsoft Systems Center Configuration Manager (SCCM)

You can use SCCM to perform distributed installations of DB2 products. SCCM deploys operating systems, software applications, and software updates.

To install a DB2 product by using SCCM, your system must be an English version of a Windows operating system, and the authorization ID for the LocalSystem account must be SYSTEM.

You can set up an installation from a central location using SCCM. This installation method is ideal if you want to roll out an installation based on the same setup on a large number of clients.

For details on how to use SCCM for distributed installation of DB2 products, see the SCCM documentation at <http://technet.microsoft.com/en-us/library/bb735860.aspx>.

If you were using Microsoft Systems Management Server (SMS), see the Getting started documentation for SCCM <http://technet.microsoft.com/en-us/library/bb694263.aspx> that contains details about upgrading from SMS to SCCM.

Response file error codes (Windows)

During installation and uninstallation on Windows platforms using response files, you may encounter error code values.

The following tables describe error return codes (primary and secondary) that can be encountered during a response file installation or uninstallation.

Table 49. Primary response file installation error codes

Error code value	Description
0	The action completed successfully.
1	The action returns a warning.
1603	A fatal error occurred.
3010	The installation or uninstallation is successful, however a reboot is required to complete the process. This does not include installations where the ForceReboot action is run. This error code is not available on Windows Installer version 1.0.

Table 50. Secondary response file installation error codes

Error code value	Description
3	The path was not found.
5	Access was denied.
10	An environment error occurred.
13	The data is invalid.
87	One of the parameters was invalid.
1602	The installation was cancelled by the user.
1610	The configuration data is corrupt. Contact your support personnel.
1612	The installation source for this product is not available. Verify that the source exists and that you can access it.
1618	Another installation is already in progress. Complete that installation first before proceeding with this installation.
1622	There was an error opening the installation log file. Verify that the specified log file location exists and that it is writable.
1632	The Temp folder is either full or inaccessible. Verify that the Temp folder exists and that you can write to it.
1633	This installation package is not supported on this platform.
1638	Another version of this product is already installed. Installation of this version cannot continue.
1639	Invalid command line argument.

For more information regarding response file return codes, refer to the Microsoft Web site.

Response file installation using a batch file (Windows)

You can use a batch file to start a response file installation.

Procedure

To start a response file installation using a batch file:

1. Edit or create your response file.
2. Create your batch file using a text editor. For example, create a batch file called `ese.bat` with the following content to install DB2 Enterprise Server Edition:

```
c:\db2ese\setup /U c:\PROD_ESE.rsp
echo %ERRORLEVEL%
```

where /U specifies the location of the response file and echo %ERRORLEVEL% specifies that you want the batch process to display the installation return code.

3. Run the batch file, using the **ese.bat** command issued from a command prompt.

Uninstalling a DB2 product, feature, or language using a response file (Windows)

These instructions guide you through the process of uninstalling a DB2 product, feature, or language using a response file on a Windows platform.

Before you begin

Before you begin the uninstallation, ensure that you have all the required user accounts to perform the uninstallation.

About this task

To silently uninstall DB2 products, features, or languages in a DB2 copy, use the **db2unins** command with the **-u** parameter.

If you have DB2 products installed in various installation paths, you must run this command separately from each installation path. Sample response file, **db2un.rsp**, is in *DB2DIR\install*, where *DB2DIR* specifies the path where the DB2 product was installed.

If you have a clustered environment, before uninstalling your DB2 product you must run the **db2mscs** command, with the **-u** parameter, from the same server that originally ran the **db2mscs** command to create the failover infrastructure. For details, see the **db2mscs** command.

Procedure

To perform an uninstallation:

1. Customize the **db2un.rsp** response file. To activate an item in the response file, remove the asterisk (*) to the left of the keyword. Then, replace the current setting to the right of the value with the new setting. The possible settings are listed to the right of the equal sign.
2. Run the **db2unins** command. For example, **db2unins -u c:\db2un.rsp**. The **db2unins** command is found under the *SQLLIB\BIN* directory.
3. Check the messages in the log file when the uninstallation finishes. The default log file location is *My Documents\DB2LOG\db2un_timestamp.log*.

Chapter 33. The response file generator

You can use the response file generator to re-create an exact setup on other machines.

The response file generator creates a response file from an existing installed and configured DB2 database product. You can use the generated response file to install an identical product setup on other machines.

For example, to connect to various databases across your network, you can install and configure an IBM Data Server Client. After this Data Server Client is installed and configured to access all databases your users have access to, run the response file generator to create a response file and a configuration profile for each DB2 instance.

Run the response file generator with the **db2rspgn** command. A response file is created for the installation and instance profiles for each instance that you specify. The response file name is *db2product_abbreviation.rsp*, for example *db2ese.rsp*. The instance profile file name is *instance_name.ins*, for example *db2inst1.ins*. You can use the response file generator to create multiple response files, one for each installed product. For example, if DB2 Enterprise Server Edition and IBM Data Server Client are installed in the same location, the **db2rspgn** command generates response files *db2ese.rsp* and *db2client.rsp*.

Note: If the instance is a DB2 pureScale instance, it is not added to the generated response file and the instance configuration file is not created.

Although several DB2 copies might already be installed on your system, the response file generator generates the response file only for the current copy. The current copy is the DB2 copy where you ran the **db2rspgn** command.

Chapter 34. Available sample response files

You can use the sample response files to install or uninstall DB2 products, features, and languages. The DB2 DVD includes ready-to-use sample response files with default entries.

The sample response files are located in:

`db2/platform/samples`

where *platform* refers to the appropriate hardware platform.

On Windows, the sample response files for IBM Data Server Runtime Client and IBM Data Server Driver Package are located in:

`\samples`

After an install, the uninstall sample response file is also saved in the `DB2DIR/install` directory, where `DB2DIR` is the full path name where you install the DB2 product.

Note: On Windows, an uninstall sample response file does not exist for IBM Data Server Runtime Client or IBM Data Server Driver Package, because the **db2unins** command is not supported for this client or driver.

Chapter 35. Sample response files (Linux, UNIX, and Windows)

You can use the sample response files to install or uninstall DB2 for Linux, UNIX, and Windows products, features, and languages. The DB2 installation DVD includes ready-to-use sample response files with default entries.

The sample response files are located in the `db2/platform/samples` directory, where *platform* refers to the appropriate hardware platform.

Table 51. Sample response files for installation of DB2 products on Linux and UNIX operating systems

Action	DB2 product	Sample response file location	Link
Install	Data Server Client	db2/platform/samples/ db2client.rsp	db2client
Install	DB2 Enterprise Server Edition	db2/platform/samples/ db2ese.rsp	db2ese
Install	DB2 Express Edition	db2/platform/samples/ db2exp.rsp	db2exp
Install	Data Server Runtime Client	db2/platform/samples/ db2rtcl.rsp	db2rtcl
Install	DB2 Workgroup Server Edition for Linux, UNIX, and Windows	db2/platform/samples/ db2wse.rsp	db2wse
Uninstall	All products	db2/platform/samples/ db2un.rsp	db2un

Table 52. Sample response files for installation of DB2 products on Windows operating systems

Action	DB2 product	Sample response file location	Link
Install	Data Server Client	db2/platform/ samples/ db2client.rsp	db2client
Install	Data Server Client	db2/platform/ samples/db2ese.rsp	db2ese
Install	DB2 Express Edition	db2/platform/ samples/db2exp.rsp	db2exp
Install	DB2 Fix Pack	db2/platform/ samples/ db2fixpk.rsp	db2fixpk
Install	DB2 Workgroup Server Edition for Linux, UNIX, and Windows	db2/platform/ samples/db2wse.rsp	db2wse
Uninstall	All products	db2/platform/ samples/db2un.rsp	db2un

Chapter 36. Response file keywords

You can use response file keywords to install or uninstall DB2 database products or features.

Response files use keywords. You can use a response file to perform tasks such as: Sample response files are provided for both an installation and uninstallation. The following response file keywords are explained in a sample response file. The sample response file can be edited. The edited response file must then be copied to your shared network drive or network file system where it will be used by your installation server.

You can use the following response file keywords to install a DB2 database product or feature, or language. On Linux or UNIX operating systems, some keywords are valid for root installation only.

PROD Specifies the product that you want to install. You can install only one DB2 database product at a time. Therefore, you can specify only one product with this keyword.

Note:

1. You should not comment out the **PROD** keyword as you might have some missing components even with a successful response file installation.
2. You cannot use the **PROD** keyword with the **db2isetup** command **-r response_file** parameter.

FILE Specifies the destination directory for a DB2 database product.

Windows

On Windows operating systems, this keyword is limited to 110 characters.

Linux or UNIX

On Linux or UNIX operating systems, this keyword is mandatory for root installations.

DB2 installation paths have the following rules:

- Can include lowercase letters (a-z), uppercase letters (A-Z), and the underscore character (_)
- Cannot exceed 128 characters
- Cannot contain spaces
- Cannot contain non-English characters

INSTALL_OPTION

Windows operating systems only and specific to DB2 Enterprise Server Edition. This keyword specifies the installation. The default is SINGLE_PARTITION.

The options are:

- INSTANCE_OWNING
- NEW_NODE
- SINGLE_PARTITION

LIC_AGREEMENT

Indicates that you have read and agree to the license agreement file in the `db2/license` directory on the DB2 database product DVD. The default is `DECLINE`.

The options are:

- `DECLINE`
- `ACCEPT`

Note:

1. You must change this keyword to `ACCEPT` to signify your acceptance of the DB2 license agreement for the installation to proceed.
2. You cannot use this keyword with the **db2isetup** command `-r response_file` parameter.

INTERACTIVE

Linux or UNIX operating systems only. Determines if installation provides prompting for location of the national language package (either on DVDs or file set location) and progress reports. The default is `NONE`.

The options are:

- `NONE`
- `YES`
- `MACHINE`

Note:

1. When `NONE` is specified, there is no interaction.
2. When `YES` is specified, progress information and prompt information is displayed to the console. During installation, if a language is missing, you are prompted for the location of the national language package.
3. Specify `MACHINE` to receive progress or prompt information in a format that is easily parsed by another program. If you are embedding the DB2 installation image, a sample program is provided.

CONFIG_ONLY

This keyword is deprecated and will be discontinued in a future release. Linux or UNIX operating systems only. Specifies that this response file is for performing configuration tasks only. For example, create an instance by running the **db2isetup** command. The default is `NO`.

The options are:

- `YES`
- `NO`

DB2SDRIVER_CFG_SOURCE_PATH

Windows operating systems only. Starting with DB2 Version 10.1 Fix Pack 2 and later fix packs, specifies the full path of a customized `db2dsdriver` configuration file. The file is copied to the common application data path during silent installation as `db2dsdriver.cfg`.

DB2_ODBC_DSN_TYPE

Windows operating systems only. Starting with DB2 Version 10.1 Fix Pack 2 and later fix packs, specifies the type of an ODBC data source name (DSN) during silent installation.

The options are:

- **USER:** Installation process registers DSNs that are listed in the `db2dsdriver.cfg` file as user DSNs for ODBC.
- **SYSTEM:** Installation process registers DSNs that are listed in the `db2dsdriver.cfg` file as system DSNs for ODBC.
- **BOTH:** Installation process registers DSNs that are listed in the `db2dsdriver.cfg` file as both user and system DSNs for ODBC.

DB2_ODBC_DSN_ACTION

Windows operating systems only. Starting with DB2 Version 10.1 Fix Pack 2 and later fix packs, specifies whether to add ODBC DSNs or clear stale ODBC DSNs during silent installation.

The options are:

- **ADD:** Installation process adds the available DSNs from the `db2dsdriver.cfg` file to be used as ODBC DSNs.
- **REPLACE:** Installation process deletes all the DSNs for the current copy before adding the available DSNs from the `db2dsdriver.cfg` file to be used as ODBC DSNs.

INSTALL_TYPE

Specifies the type of installation.

The options are:

- **COMPACT**
- **TYPICAL**
- **CUSTOM**

A compact or typical installation type ignores any custom keywords (**COMP**).

Note: You cannot use this keyword with the **db2isetup** command **-r** *response_file* parameter.

DB2_COPY_NAME

Windows operating systems only. An identifier that represents the name that is used to refer to a set of DB2 database products that are installed in the same location. This name is limited to 64 characters.

DEFAULT_COPY

Windows operating systems only. Specifies if the copy that is being installed or modified will be set as the default DB2 copy used by DB2 applications. The default is **N0** unless it is the only DB2 copy installed.

The options are:

- **N0**
- **YES**

COPY_NAME

Windows operating systems only. Represents the name that is used to refer to the IBM Data Server Driver Package installed. Valid only when the **PROD** keyword is set to **IBM_DATA_SERVER_DRIVER**. This name is limited to 64 characters.

DEFAULT_CLIENT_INTERFACE_COPY

Windows operating systems only. Specifies if the copy of DB2 database products, or the IBM Data Server Driver Package, installed will be set as the default IBM database client interface copy. If this keyword is set to **YES**, the IBM data server client interface (ODBC/CLI driver and .NET data provider) in the copy is the default driver to be used by the applications. If you are installing the IBM Data Server Driver, the default is **N0** unless the

computer does not have a DB2 database product, or the IBM Data Server Driver Package installed. If you are installing other products, the default is the value of the **DEFAULT_COPY** keyword.

The options are:

- YES
- NO

AUTOSTART_SSH_SERVER

Windows operating systems only. Specifies whether you want to autostart the IBM SSH Server for Windows service. The default is to autostart the service.

The options are:

- YES
- NO

SSH_SERVER_INSTALL_DIR

Windows operating systems only. Determines the base installation path for the IBM Secure Shell (SSH) Server for Windows service. This keyword is only evaluated when the SSH Server is first installed. If this keyword is not specified in the response file and SSH is to be installed, the default installation directory is used. The default directory is: %program files%\IBM\IBM SSH Server. If the default directory is not empty, the directory name is append with "_01", "_02", and so on. For example, %program files%\IBM\IBM SSH Server_01

POPULATE_XML_FROM_CATALOG

Specifies whether the db2dsdriver.cfg configuration file is populated with catalog information. The default is FALSE.

The options are:

- TRUE
- FALSE

INSTALL_TSAMP

Important: The INSTALL_TSAMP response file keyword has been deprecated in Version 10.1 and might be removed in a future release. Starting in Version 10.1, IBM Tivoli System Automation for Multiplatforms (SA MP) is automatically installed when SA MP is a required component for other components selected for installation. For more information, see “IBM Tivoli System Automation for Multiplatforms (SA MP) is now automatically installed” in *What’s New for DB2 Version 10.1*.

Specifies support of IBM Tivoli System Automation for Multiplatforms (SA MP) on AIX and Linux operating systems only. The default value is YES.

The options are:

- YES
- NO

If INSTALL_TSAMP=NO is specified, but another component for installation requires SA MP, SA MP is automatically installed regardless of the value that is specified for this keyword.

MIGRATE_PRIOR_VERSIONS

This keyword is deprecated, use the **UPGRADE_PRIOR_VERSIONS** keyword instead. Windows operating systems only. Specifies if an existing version of the DB2 database product is to be updated. The default is FALSE.

The options are:

- TRUE
- FALSE

UPGRADE_PRIOR_VERSIONS

Specifies if an existing version of the DB2 database product is to be upgraded. The default is FALSE.

The options are:

- TRUE
- FALSE

This keyword is supported on Linux, UNIX, and Windows operating systems. However, on Linux and UNIX operating systems, this keyword is valid only for non-root installation. On Windows operating systems, in addition to setting **UPGRADE_PRIOR_VERSIONS** to the value TRUE, you must also set the **DB2_COPY_NAME** keyword to the value of an existing copy name.

UPGRADE_DBCK_IGNORE_TYPE1

Linux and UNIX operating systems only, and specific to non-root installation. Forces the **db2setup** command to ignore type-1 index checking. The default is NO.

The options are:

- YES
- NO

PACKAGE_LOCATION

Linux or UNIX operating systems only. This keyword is used only when "nlpack" is removed from the image and NL languages are selected to install. This keyword specifies the location of the nlpack. The path can be the location of the product DVD or the location where the contents of the DVD are saved.

Built-in search paths allow automatic searching of the same parent directory. For example, if the contents of a DB2 database product DVD are copied into the following subdirectories:

```
/db2images/ese/dvd  
/db2images/nlpack/dvd
```

The subdirectories are automatically searched without specifying each subdirectory (including searching subdirectories for the national language DVD):

```
PACKAGE_LOCATION=/db2images/ese/disk1
```

If the contents of a DB2 database product DVD are copied into different parent directories, you can repeat this keyword as many times as required. For example:

```
PACKAGE_LOCATION=/db2images1/dvd  
PACKAGE_LOCATION=/db2images2/nlpack/dvd
```

COMP

Specifies the components that you want to install. The setup program automatically installs components that are required for a product.

In a custom installation, you must select components individually. To select components, uncomment the **COMP** keywords for the components that you want installed (components differ depending on the product).

Note: This keyword is ignored unless your **INSTALL_TYPE** is CUSTOM.

LANG Specifies which languages to install. You must uncomment any additional languages that you want to install. The English language is mandatory and is always selected.

CLIENT_IMPORT_PROFILE

Specifies a configuration profile exported by the **db2cfexp** command. The profile contains instance-related parameters.

- Windows operating systems:
 - If the value provided is a full path name, it must be on a local drive (not a remote network drive). If it is not a full path name, the configuration profile must be in the same directory as the response file.
- Linux or UNIX operating systems:
 - Specifies a file name with a full path to the configuration profile.

REBOOT Windows operating systems only. Specifies whether to restart the system when the installation or uninstallation has completed. The default is NO.

The options are:

- YES
- NO

KILL_PROCESSES

Windows operating systems only. If you have an existing version of DB2 running and this keyword is set to YES, it will terminate your running DB2 processes without a prompt.

The options are:

- YES
- NO

NEW_CONTACT

Specifies if the email address contact on the notification and contact list is new or existing. A new contact is not required if there is an existing email address. On Linux or UNIX operating systems, this keyword is valid only for root installation.

The options are:

- YES
- NO

This keyword is optional. If the keyword is not specified, by default the contact is considered as a new contact.

DB2_OLEDB_GUID

Windows operating systems only. Use this keyword to enter the GUID used by DB2 applications using OLE DB. If you do not specify this keyword, it is generated.

CREATE_DAS

Important: The DB2 Administration Server (DAS) has been deprecated in Version 9.7 and might be removed in a future release. The DAS is not supported in DB2 pureScale environments. Use software programs that use the Secure Shell protocol for remote administration. For more information, see “ DB2 administration server (DAS) has been deprecated” at .
Windows operating systems only. Specifies whether to create the Database Administration Server (DAS) or not. There can be one DAS only for all

DB2 copies on the same system. Keep this keyword commented out in the response file because if this keyword is specified and set to YES, and a DAS instance exists on the system, the installation fails.

When this keyword is commented out in the response file (or is not specified), the DB2 installer determines whether a DAS must be created. In the case where:

- A DAS instance does not exist, a DAS is created.
- A DAS instance exists, this keyword is ignored.

The options are:

- YES
- NO

CHECK_UNCPATH

Windows operating systems only. Specifies to check that the Universal Naming Convention (UNC) path is in the **PATH** environment variable. The default value is YES.

The options are:

- YES
- NO

DB2_COMMON_APP_DATA_TOP_PATH

Windows operating systems only. Specifies a user-defined location to hold the DB2 common application data for the DB2 copy to be installed. You can use this keyword only during the initial installation of DB2 product in a new location. The path specified for this keyword is used to populate the **DB2_COMMON_APP_DATA_PATH** registry variable. DB2 common application data is placed in the IBM\DB2\<COPYNAME> subdirectory, which is created under a path specified in the **DB2_COMMON_APP_DATA_PATH** registry variable.

For example, when the DB2 product is installed with the **DB2_COMMON_APP_DATA_TOP_PATH** keyword set to a specific path:

DB2_COMMON_APP_DATA_TOP_PATH=C:\User\MyCommonAppData\

After the DB2 product is successfully installed, the **DB2_COMMON_APP_DATA_PATH** registry variable is set to following value:

DB2_COMMON_APP_DATA_PATH=C:\User\MyCommonAppData\

DB2 common application data is stored in following path :

C:\User\MyCommonAppData\IBM\DB2\DB2COPY1\

After installation, this variable cannot be altered.

If this keyword is not set, the Windows' default common application data path is used to hold the DB2 common application data.

- For Windows XP and Windows 2003 operating systems: C:\Documents and Settings\All Users\Application Data\
- For Windows Vista and later operating systems: C:\ProgramData\

Note:

1. You cannot use this keyword to change the DB2 Common Application Data path during a DB2 upgrade.
2. Use of the network drive is not supported.
3. You can use only alphanumeric characters for the **DB2_COMMON_APP_DATA_TOP_PATH** path.

DB2 Administration Server settings

To enable any of the following DB2 Administrator Server (DAS) settings, remove the * (uncomment). This setting is applicable for Windows, Linux, and UNIX environments but on Linux or UNIX operating systems, this keyword is valid for root installation only.

- On Linux or UNIX operating systems:
 - *DAS_USERNAME = dasuser
 - *DAS_PASSWORD = dasp
 - *DAS_GID = 100
 - *DAS_UID = 100
 - *DAS_GROUP_NAME = dasgroup
 - *DAS_SMTP_SERVER = jsmith.torolab.ibm.com
- On Windows operating systems:
 - *DAS_USERNAME = dasuser
 - *DAS_DOMAIN = domain
 - *DAS_PASSWORD = dasp
 - *DAS_SMTP_SERVER = jsmith.torolab.ibm.com

You can also specify LOCALSYSTEM as a valid user name that does not require a password. The LOCALSYSTEM user name is available for all products, except DB2 Enterprise Server Edition.

The following options specify where the DAS contact list will be kept. On Linux or UNIX operating systems, these options are valid for root installation only. If the contact list is remote, you must specify a user name and password that has authority to add a contact to the system.

```
*DAS_CONTACT_LIST = LOCAL or REMOTE (DEFAULT = LOCAL)
*DAS_CONTACT_LIST_HOSTNAME = hostname
*DAS_CONTACT_LIST_USERNAME = username
*DAS_CONTACT_LIST_PASSWORD = password
```

Special instance specifications

The special instance specifications takes the instance section not instance name. The instance section must exist in the response file.

- Windows operating systems:
 - **DEFAULT_INSTANCE** - This is the default instance.
- Linux or UNIX operating systems:
 - none

Instance specifications

For root installations, you can use the response file to create as many instances as you want. To create an instance you must specify an instance section with the **INSTANCE** keyword. Any keywords that contain the value specified in **INSTANCE** as a prefix belong to that instance. On Windows operating systems, you can also specify LOCALSYSTEM as a valid user ID that does not require a password. The LOCALSYSTEM user ID is available for all products, except DB2 Enterprise Server Edition.

The **START_DURING_INSTALL** keyword determines if the instance is started after installation. The **AUTOSTART** keyword determines if the instance is auto started after system reboot. If **START_DURING_INSTALL** is not specified, **AUTOSTART** also determines if the instance is started after installation.

The following are examples of instance specifications for Windows, Linux, and UNIX operating systems:

To create an instance for an existing user, do not specify the *PASSWORD*, *UID*, *GID*, *GROUP_NAME*, or *HOME_DIRECTORY* keywords. To create the user or the fenced user, you must specify all of the previous keywords in the **db2setup** command.

- On Linux or UNIX operating systems:

```
*INSTANCE=DB2_INST
*DB2_INST.NAME = db2inst1
*DB2_INST.TYPE = ESE
*DB2_INST.PASSWORD = PASSWORD
    (valid for root installation only)
*DB2_INST.UID = 100
    (valid for root installation only)
*DB2_INST.GID = 100
    (valid for root installation only)
*DB2_INST.GROUP_NAME = db2grp1
    (valid for root installation only))
*DB2_INST.HOME_DIRECTORY = /home/db2inst1
    (valid for root installation only)
*DB2_INST.SVCENAME = db2cdb2inst1
*DB2_INST.PORT_NUMBER = 50000
*DB2_INST.FCM_PORT_NUMBER = 60000
    (valid for root installation only)
*DB2_INST.MAX_LOGICAL_NODES = 4
    (valid for root installation only)
*DB2_INST.AUTOSTART = YES
*DB2_INST.START_DURING_INSTALL = YES
*DB2_INST.FENCED_USERNAME = USERNAME
    (valid for root installation only)
*DB2_INST.FENCED_PASSWORD = PASSWORD
*DB2_INST.FENCED_UID = 100
*DB2_INST.FENCED_GID = 100
*DB2_INST.FENCED_GROUP_NAME = db2grp1
*DB2_INST.FENCED_HOME_DIRECTORY = /home/db2inst1
*DB2_INST.CONFIGURE_TEXT_SEARCH = NO
*DB2_INST.TEXT_SEARCH_HTTP_SERVICE_NAME = db2j_DB2_INSTANCE
    (valid for root installation only)
*DB2_INST.TEXT_SEARCH_HTTP_PORT_NUMBER = 55000*DB2_INST.FEDERATED = YES
*DB2_INST.INFORMIX_CLIENT_HOME = BLANK
    (valid for root installation only)
*DB2_INST.INFORMIX_SERVER_HOME = BLANK
```

- To create multiple instances:

```
*INSTANCE=DB2_INS2
*DB2_INS2.NAME = db2inst2
*DB2_INS2.TYPE = ESE
*DB2_INS2.PASSWORD = PASSWORD
    (valid for root installation only)
*DB2_INS2.UID = 100
    (valid for root installation only)
*DB2_INS2.GID = 100
    (valid for root installation only)
*DB2_INS2.GROUP_NAME = db2grp1
    (valid for root installation only))
*DB2_INS2.HOME_DIRECTORY = /home/db2inst1
    (valid for root installation only)
*DB2_INS2.SVCENAME = db2cdb2inst1
*DB2_INS2.PORT_NUMBER = 50000
*DB2_INS2.FCM_PORT_NUMBER = 60000
    (valid for root installation only)
*DB2_INS2.MAX_LOGICAL_NODES = 4
    (valid for root installation only)
*DB2_INS2.AUTOSTART = YES
*DB2_INS2.START_DURING_INSTALL = YES
*DB2_INS2.FENCED_USERNAME = USERNAME
    (valid for root installation only)
*DB2_INS2.FENCED_PASSWORD = PASSWORD
```

```

*DB2_INS2.FENCED_UID = 100
*DB2_INS2.FENCED_GID = 100
*DB2_INS2.FENCED_GROUP_NAME = db2grp1
*DB2_INS2.FENCED_HOME_DIRECTORY = /home/db2inst1
*DB2_INS2.CONFIGURE_TEXT_SEARCH = NO
*DB2_INS2.TEXT_SEARCH_HTTP_SERVICE_NAME = db2j_DB2_INSTANCE
    (valid for root installation only)
*DB2_INS2.TEXT_SEARCH_HTTP_PORT_NUMBER = 55000*DB2_INST.FEDERATED = YES
*DB2_INST.INFORMIX_CLIENT_HOME = BLANK
    (valid for root installation only)
*DB2_INST.INFORMIX_SERVER_HOME = BLANK

```

- On Windows operating system:

```

DEFAULT_INSTANCE = DB2_INST
INSTANCE = DB2_INST
DB2_INST.NAME = DB2
*DB2_INST.TYPE = ESE
DB2_INST.USERNAME = username
*DB2_INST.DOMAIN = domainname
DB2_INST.PASSWORD = password
DB2_INST.AUTOSTART = YES
DB2_INST.START_DURING_INSTALL = YES
*DB2_INST.MAX_LOGICAL_NODES = 4
*DB2_INST.CONFIGURE_TEXT_SEARCH = NO
*DB2_INST.TEXT_SEARCH_HTTP_SERVICE_NAME = db2j_DB2_INSTANCE
*DB2_INST.TEXT_SEARCH_HTTP_PORT_NUMBER = 55000

```

- To create multiple instances:

```

*INSTANCE = DB2_INS2
*DB2_INS2.NAME = DB2_01
*DB2_INS2.TYPE = ESE
*DB2_INS2.USERNAME = username
*DB2_INS2.DOMAIN = domainname
*DB2_INS2.PASSWORD = password
*DB2_INS2.AUTOSTART = YES
*DB2_INS2.START_DURING_INSTALL = YES
*DB2_INS2.MAX_LOGICAL_NODES = 4
*DB2_INS2.CONFIGURE_TEXT_SEARCH = NO
*DB2_INS2.TEXT_SEARCH_HTTP_SERVICE_NAME = db2j_DB2_INSTANCE
*DB2_INS2.TEXT_SEARCH_HTTP_PORT_NUMBER = 55000

```

Note: On Windows operating system, in the case of the second and subsequent instances, the DEFAULT_INSTANCE keyword is not specified. This keyword is specified on the first instance only.

Database Section

These keywords can be used to have the installation create or catalog a database on the machine that is being installed.

```

DATABASE = DATABASE_SECTION
DATABASE_SECTION.INSTANCE = db2inst1
DATABASE_SECTION.DATABASE_NAME = MYDB
DATABASE_SECTION.LOCATION = LOCAL
DATABASE_SECTION.ALIAS = MYDB
DATABASE_SECTION.USERNAME = username
DATABASE_SECTION.PASSWORD = password

```

* these keywords are only used for REMOTE databases that are being cataloged

```

DATABASE_SECTION.SYSTEM_NAME = hostname
DATABASE_SECTION.SVCENAME = db2c_db2inst1

```

Contact Section

These keywords define a contact section that will be created by the installation process if it does not yet exist. The Health notifications for the instance that is specified will be sent to this contact.


```

CONTACT = contact_section
contact_section.NEW_CONTACT = YES
contact_section.CONTACT_NAME = contact name
contact_section.INSTANCE = DB2_INSTANCE
contact_section.EMAIL = Email address
contact_section.PAGER = NO

```

On Linux or UNIX operating systems, these keywords are valid for root installation only.

DB2 Information Center settings section

The default location for accessing the DB2 documentation is the IBM website. Edit this section only if you want to access DB2 documentation from a different location, such as your local computer or an intranet server.

The following options specify the host name and port number where the Information Center server is installed, and the service name and port number that the Information Center server service uses.

```

*DB2_DOCHOST = hostname
*DB2_DOCPORT = 1024-65535
*DB2_ECLIPSEIC_SVCENAME = db2icv95
*DB2_ECLIPSEIC_PORT = 51000

```

DB2_ECLIPSEIC_SVCENAME and **DB2_ECLIPSEIC_PORT** are only valid for the installation of the *DB2 Information Center*.

Extended security section

Windows operating systems only. These keywords define a security section.

```

*DB2_EXTSECURITY           = YES
*DB2_ADMINGROUP_NAME      = DB2ADMNS
*DB2_USERSGROUP_NAME      = DB2USER
*DB2_ADMINGROUP_DOMAIN    = BLANK
*DB2_USERSGROUP_DOMAIN    = BLANK

```

No configuration option

Windows operating systems only. This keyword gives you the option to install a DB2 database product with only the mandatory configuration. The DB2 instance must be created manually. The default is NO.

```
*NO_CONFIG      = NO
```

The options are:

- YES
- NO

IBM DB2 pureScale Feature installation keywords

The following response file keywords are used in a DB2 pureScale Feature installation. Special considerations and interdependencies are required when using these keywords.

HOST Specifies the hosts to which the DB2 product is to be installed. There can be multiple occurrences of this HOST keyword in the response file, one of each distinct host. This keyword is limited to 14 characters, and is a mandatory keyword for a DB2 pureScale instance or update. The maximum number of hosts is 130 (two cluster caching facilities, and 128 members.) This keyword is used as the ID of a host set, to differentiate the other host sets, and acts as the prefix of the keyword HOSTNAME and CLUSTER_INTERCONNECT_NETNAME. For example:

```

HOST = host1
host1.HOSTNAME = hostname1
host1.CLUSTER_INTERCONNECT_NETNAME = ibname1

HOST = host2
host2.HOSTNAME = hostname2
host2.CLUSTER_INTERCONNECT_NETNAME = ibname2

```

HOSTNAME

Specifies a host name to a physical machine or LPAR on which a DB2 member or cluster caching facility is running. This host name must be a unique host name across the network, and is limited to 255 characters. This keyword is prefixed with the HOST keyword value, and is a mandatory keyword for a DB2 pureScale instance or update. If an alias or IP is provided, the name is converted to a canonical host name for instance creation. The host name format can be: long name (for example, hostname1.ibm.com), short name (for example, hostname1), or an IPv4 or IPv6 address.

CLUSTER_INTERCONNECT_NETNAME

Specifies the IP address or host name of the interconnect used for high speed communication between members and cluster caching facilities in the DB2 instance. This keyword is limited to 255 characters, and is a mandatory keyword for a DB2 pureScale instance or update. The host high speed name format can be: long name (for example, host1.ibm.com), short name (for example, host1), or an IPv4 or IPv6 address. You must use the same format for all hosts specified.

DB2_CLUSTER_SERVICES_TIEBREAKER_DEVICE_PATH

Identifies the device path being used to set up a cluster services tiebreaker. All hosts specified in the HOST sections must have direct access to this path. A shared disk is required for cluster recovery purposes. This disk should be at least 25 MB. For example, /dev/hdisk2. However, this keyword cannot have the same value as keyword

INSTANCE_SHARED_DEVICE_PATH.

INSTANCE_SHARED_DEVICE_PATH

Identifies the device path used to create a DB2 Managed clustered file system for shared files from a DB2 pureScale instance. All hosts specified in the HOST sections must have direct access to this path. This path must be at least 2 GBs. This is a mandatory keyword for the first DB2 pureScale instance. For example, /dev/hdisk3.

INSTANCE_SHARED_MOUNT

Specifies the mount point for a new General Parallel File System (GPFS). The specified path must be a new and empty path that is not nested inside an existing GPFS file system.

INSTANCE_SHARED_DIR

Identifies the path to an existing file system for storing shared files from a DB2 pureScale instance. When this keyword is specified, a User-Managed file system method is used, and the value specified must be an existing GPFS type file system which is mounted on all hosts. All hosts specified in the HOST sections must have direct access to this directory. This partition must be at least 1 GB. If either the **INSTANCE_SHARED_DEVICE_PATH** keyword is specified, or a DEFAULT_INSTPROF record exists in the Global Registry, this keyword is ignored.

Instance specifications

The following are examples of instance specifications for the DB2 pureScale Feature:

- On AIX operating systems:

```
INSTANCE = db2inst1
db2inst1.NAME = db2sdin1
*db2inst1.UID = 100
db2inst1.GROUP_NAME = db2isdml
*db2inst1.GID = 100
db2inst1.HOME_DIRECTORY = /home/db2inst1
db2inst1.PASSWORD = password
db2inst1.START_DURING_INSTALL = YES

db2inst1.PREFERRED_PRIMARY_CF = host1
db2inst1.PREFERRED_SECONDARY_CF = host3
db2inst1.MEMBER = host2
db2inst1.MEMBER = host4

*db2inst1.SVCENAME = db2c_db2sdin1
*db2inst1.PORT_NUMBER = 50000
*db2inst1.FCM_PORT_NUMBER = 60000
*db2inst1.MAX_LOGICAL_NODES = 1
*db2inst1.DB2CF_PORT_NUMBER = 56001
*db2inst1.DB2CF_MGMT_PORT_NUMBER = 56000
*db2inst1.TYPE = DSF
*db2inst1.AUTHENTICATION = SERVER
```

Note:

- The keyword **PREFERRED_PRIMARY_CF** identifies the host that is designated as the preferred primary cluster caching facility for the instance being set up. This is a mandatory keyword for a DB2 pureScale instance.
- The keyword **PREFERRED_SECONDARY_CF** identifies the host that is designated as the secondary cluster caching facility for the instance being set up.
- The keyword **MEMBER** identifies the host that is designated as a member for the instance being set up. This is a mandatory keyword for a DB2 pureScale instance.
- The keywords **DB2CF_PORT_NUMBER** and **DB2CF_MGMT_PORT_NUMBER** are reserved for the cluster caching facility. The cluster caching facility is an instance designation, therefore the port names will be predefined based on the name of the instance. For example: When the instance name is db2sdin1, the names required for the cluster caching facility ports are DB2CF_db2sdin1 and DB2CF_db2sdin1_MGMT_PORT_NUMBER. Ensure that all port names are the same on all participating hosts. The two cluster caching facility ports must be available on all the participating hosts. The DB2CF port keywords are optional. If they are not specified, the DB2 installer generates unique ports across all hosts.

Keywords for the db2rfe configuration file

Use the following keywords to select features and settings to create the **db2rfe** configuration file:

INSTANCENAME

Specifies the user name who owns the non-root installation copy.

Note: Do not comment out the **INSTANCENAME** keyword. This is a required keyword.

SET_ULIMIT

AIX only. For other operating systems, see the system documentation for information about setting this manually. On AIX, when this keyword is set to YES, the **db2rfe** command sets the hard and soft data parameter to unlimited, the hard and soft file size to unlimited, and sets the hard and soft nofile parameter to 65536.

The options are:

- YES
- NO

The default is NO.

ENABLE_DB2_ACS

AIX, Linux on AMD64/EM64T, and Linux on Power Systems™ servers only. Enables the DB2 Advanced Copy Services (ACS). The default value is NO

The options are:

- NO
- YES

ENABLE_HA

Specifies support of High Availability using IBM Tivoli System Automation for Multiplatforms. The default value is NO.

The options are:

- NO
- YES

ENABLE_OS_AUTHENTICATION

Specifies support of operating system authentication for database connection. The default value is NO.

The options are:

- NO
- YES

RESERVE_REMOTE_CONNECTION

Reserves service for a remote connection. The default value is NO.

The options are:

- NO
- YES

Note:

1. If the database manager configuration parameter **svcename** has a value, and the keywords **SVCENAME** or **SVCEPORT** are also provided in the configuration file, all three values must match.
2. If the database manager configuration parameter **svcename** has a value, and either keyword **SVCENAME** or **SVCEPORT** is not provided, either **port_number** or **service_name** of **svcename** is used with a generated **service_name** or **port_number** respectively.
3. If the database manager configuration parameter **svcename** is not set, and neither keyword **SVCENAME** and **SVCEPORT** are provided, a default service name and port number is generated and reserved.

SVCENAME

Used with the keyword **RESERVE_REMOTE_CONNECTION**. Specifies the service name of a remote connection. This keyword is limited to 14 characters.

SVCEPORT

Used with the keyword **RESERVE_REMOTE_CONNECTION**. Specifies a port number of a remote connection within the range 1024 - 65535.

RESERVE_TEXT_SEARCH_CONNECTION

Reserves the service entry for the port used by DB2 Text Search. The default value is NO.

The options are:

- NO
- YES

SVCENAME_TEXT_SEARCH

Used with the keyword **RESERVE_REMOTE_CONNECTION**. Specifies the service name of the service entry for the port used by DB2 Text Search. This keyword is limited to 14 characters.

SVCEPORT_TEXT_SEARCH

Used with the keyword **RESERVE_REMOTE_CONNECTION**. Specifies a port number of the service entry for the port used by DB2 Text Search.

Uninstallation keywords

You can use the following response file keywords to uninstall a product or feature:

The following response file keywords can be used to uninstall a product, feature, or language:

REMOVE_PROD

Specifies the DB2 database product that you want to uninstall. You can uninstall DB2 database products individually or you can uninstall all DB2 database products. To uninstall all products in the current DB2 copy, specify **REMOVE_PROD=ALL**. The DB2 products in other DB2 copies are not affected. On Windows operating systems, you cannot combine this keyword with **REMOVE_COMP** or **REMOVE_LANG**.

REMOVE_COMP

Specifies the DB2 components that you want to uninstall. On Windows operating systems, you cannot combine this keyword with **REMOVE_PROD**.

REMOVE_LANG

Specifies the languages you want to uninstall. You cannot uninstall English. To uninstall all languages, except English, from the current DB2 copy, specify **REMOVE_LANG=ALL**. English will not be removed until the last DB2 product in the DB2 copy is uninstalled. On Windows operating systems, you cannot combine this keyword with **REMOVE_PROD**.

REMOVE_TSAMP

Specifies the IBM Tivoli System Automation for Multiplatforms to be uninstalled. The default is YES.

The options are:

- YES
- NO

This keyword is considered only when a DB2 copy is removed. To remove a DB2 copy, specify each product or uncomment REMOVE_PROD=ALL.

When removing DB2 pureScale Feature, if Tivoli SA MP was originally installed as part of the DB2 installation, it is also automatically uninstalled when the last DB2 pureScale copy is removed. If you want to keep the RSCT peer domain but remove the DB2 copy, you can skip the removal by specifying REMOVE_TSAMP=NO.

REMOVE_GPFS

Specifies the IBM General Parallel File System (GPFS) to be uninstalled. The default is YES.

The options are:

- YES
- NO

This keyword is only considered when a DB2 copy is removed. To remove a DB2 copy, specify each product or uncomment REMOVE_PROD=ALL.

When removing DB2 pureScale Feature, if GPFS was originally installed as part of the DB2 installation, it is also automatically uninstalled when the last DB2 pureScale copy is removed. If you want to keep the GPFS cluster and its file systems but remove the DB2 copy, you can skip the removal by specifying REMOVE_GPFS=NO.

REMOVE_DAS

Windows operating systems only. Determines if the DAS is to be deleted. The default is NO.

The options are:

- NO
- YES

Depending on your environment setup, the installation action varies.

- If you have only one DB2 copy with only one server product installed, and the server product is being removed, the DAS is removed without checking the value of this keyword.
- If you have only one DB2 copy and 2 or more server products installed, and one server product is being removed, the value of this keyword is not checked and the DAS is not removed.
- If you have two or more DB2 copies, this keyword is only examined when both the following are also true:
 - The DAS is configured and active under the DB2 copy that is being removed, and
 - You are removing all server products in the DB2 copy.

In this case, when this keyword is set to NO, and the conditions specified previously are met, the uninstallation stops and does not continue because the keyword is set to NO (do not remove DAS).

When this keyword is set to YES, and the conditions specified previously are met, the uninstallation always remove the DAS when removing the current DB2 copy. In this case, other DB2 copies that require the use of a DAS will not function properly after the DAS is removed. Before removing the DAS, you must move the DAS to another DB2 copy using the **dasupdt** command.

REMOVE_INSTALL_CREATED_USERS_GROUPS

Removes the users and groups created by the DB2 installer for the current DB2 copy. The default is NO.

The options are:

- YES
- NO

This keyword is only considered when removing all products or the final product within the DB2 copy.

REBOOT Windows operating systems only. Specifies whether to restart the system when the installation or uninstallation has completed. The default is NO.

The options are:

- YES
- NO

Chapter 37. Embedding the DB2 installation image (Linux and UNIX)

On Linux and UNIX operating systems, you can embed a DB2 installation image into the installation image of your own application.

When installing a DB2 product using a response file with the interactive keyword specified, installation-related information such as progress or prompts is provided in a format easily parsed by your application.

Procedure

To bundle a DB2 installation image in your own application:

1. Copy the DB2 sample program, located in `db2/samples/`, to a working directory. Sample programs are provided in both C and Java. The C and Java subdirectories contain a sample program and a readme file.
2. Build the sample program using the makefile provided or a compatible compiler.
3. Modify the response file to specify the keyword `INTERACTIVE=MACHINE`.
4. Using the sample program, start the DB2 installation from your working directory.
 - In a C-based installation application, enter:
`./InstallTester image -r response_file`
 - In a Java based installation application, enter:
`java InstallTester image -r response_file`

where

- *image* represents the location of the DB2 installable image where either the **db2setup** command or the **db2_install** command is located
- *response_file* specifies the full path and file name of the response file to use.

Chapter 38. Exporting and importing a profile

You can import and export configuration information to another DB2 workstation instance.

About this task

If you did not use a configuration profile when you installed your DB2 product using the response file that was created by the response file generator, you can create a configuration file and import it to another workstation.

Procedure

1. To create a configuration profile, enter the **db2cfexp** command specifying the fully qualified name of the target export file. The resulting profile contains only configuration information associated with the current DB2 database instance.
2. To import the configuration profile, you can:
 - Use the **db2cfimp** command
 - Use a response file by uncommenting the keyword `DB2.CLIENT_IMPORT_PROFILE` and specify the *filename* as the export file

Chapter 39. Stopping DB2 processes during an interactive installation (Windows)

If you are modifying or installing a new product on top of an existing DB2 copy, it is recommended that you stop all DB2 processes running on the DB2 copy before proceeding.

About this task

You must exercise extreme caution when you stop active DB2 processes because the termination of a DB2 process can cause the loss of data.

To lessen the risk of data loss, it is recommended that you issue the **db2stop** command for each instance.

The following describes how to stop DB2 processes.

Procedure

1. For an interactive installation, to stop any running DB2 processes, specify the **/F** option for the **setup** command. The **/F** option stops the running processes, and the message and prompt are not displayed.
2. In addition, DB2 services can be viewed in the Services Window to ensure that they have been stopped.

Chapter 40. Stopping DB2 processes during a response file installation (Windows)

If any DB2 processes are running when the DB2 setup command is issued, you are unable to install the DB2 product.

About this task

You must stop the DB2 processes so that the installation can proceed. You must exercise extreme caution when you stop active DB2 processes because the termination of a DB2 process can cause the loss of data. The following describes how to stop these processes.

To lessen the risk of data loss, it is recommended that you issue the **db2stop** command for each instance and that you close any applications accessing DB2 data.

For a response file installation, you can use either of the following methods to stop any active DB2 processes. If you specify either of these options, the active DB2 processes are stopped before the installation proceeds.

Procedure

- Specify the **/F** option for the **setup** command.
- Set the **KILL_PROCESSES** keyword to **YES** (the default is **NO**).

Part 6. Installing the DB2 Information Center

You can install the DB2 Information Center using the DB2 Setup wizard.

You can also install the DB2 Information Center using other options.

Chapter 41. Installing the DB2 Information Center using the DB2 Setup wizard (Linux)

Using the DB2 Setup wizard, you can define your installation preferences and install the DB2 Information Center on a computer that uses a Linux operating system.

Before you begin

Review the prerequisite information in the “Installation requirements for DB2 servers and IBM data server clients (Linux)” topic.

About this task

DB2 database product documentation can be accessed in three ways:

- at the IBM Web site
- on an intranet server
- on a version installed on your computer

By default, DB2 database products access DB2 documentation at the IBM Web site. If you want to access the DB2 documentation on an intranet server or on your own computer, you must install the documentation from the *DB2 Information Center DVD*, *Passport Advantage*, or DB2 database product documentation website.

Restrictions

- To install a network version of the *DB2 Information Center*, you must be logged on as a user with root user authority. To install a Workstation version of the DB2 Information Center, you do not require root user authority.
- You cannot install the *DB2 Information Center* in a location where a DB2 database product is installed. Similarly, the *DB2 Information Center* cannot coexist in the same installation path of a previous version of the *DB2 Information Center* on the same system.
- You can only install one copy of the *DB2 Information Center* of the same version on your system. For example, you can install a copy of the *DB2 Information Center Version 9.7* and *DB2 Information Center Version 10.1* on the same system, but you cannot install two copies of the *DB2 Information Center Version 10.1* on the same system.
- If you install the *DB2 Information Center* on a system with a firewall, and you plan for other systems to access the *DB2 Information Center*, you must open the port in your firewall setting.

Procedure

To install the *DB2 Information Center* using the DB2 Setup wizard:

1. Log on to the system.
2. If you have the *DB2 Information Center product DVD*:
 - a. Insert and mount the DVD on your system.
 - b. Change to the directory where the DVD is mounted by entering the following command:

```
cd /dvd
```

where */dvd* represents the mount point of the DVD.

3. If you downloaded the image from an online source (such as Passport Advantage, or, *ibm.com*[®]), untar the image:
 - a. Decompress the product file: `gzip -d product.tar.gz` For example,
`gzip -d ese.tar.gz`
 - b. Untar the product file: `tar -xvf product.tar` For example,
`tar -xvf ese.tar`
 - c. Change directory into the product directory: `cd product` For example,
`cd ese`
4. Enter the **./db2setup** command to start the DB2 Setup wizard.
5. The DB2 Setup Launchpad opens. From the launchpad, you can view installation prerequisites and the release notes, or you can proceed directly to the installation. You should review the installation prerequisites and release notes for late-breaking information.
6. Click **Install a Product** and the **Install a Product** window displays.
7. On the **Install a Product** page, if you do not have an existing *DB2 Information Center* installed on your computer, launch the installation by clicking **Install New**.

If you already have an existing *DB2 Information Center* installed on your computer, click **Work with Existing** to work with the existing *DB2 Information Center*.
8. On the **Welcome to the DB2 Setup wizard** page, click **Next**. The DB2 Setup wizard will guide you through the program setup process.
9. To proceed with the installation, you must accept the license agreement. On the **Software License Agreement** page, select **Accept** and click **Next**.
10. On the **Select installation, response file creation, or both** page, select **Install DB2 Information Center on this computer**. If you want to use a response file to install the *DB2 Information Center* on this or other computers at a later time, select **Install DB2 Information Center on this computer and save my settings in a response file**. You can specify where the response file will be saved. Click **Next**.
11. On the **Select the languages to install** page, select the languages the *DB2 Information Center* will install. By default, the *DB2 Information Center* is installed in the */opt/ibm/db2ic/V10.1* directory. However, you can specify your own installation path. Click **Next**.
12. Configure the *DB2 Information Center* for incoming communication on the **Specify the DB2 Information Center port** page. Click **Next** to continue the installation.

If you specify a port number other than the default and you receive the error "The service name specified is in use", you can correct this error by either choosing to use the default port number or by specifying a different service name.
13. Review the installation choices you have made in the **Start copying files** page. To change any settings, click **Back**. Click **Finish** to complete the installation of the *DB2 Information Center* files onto your computer.

Results

The installation logs, *db2setup.log* and *db2setup.err* will be located, by default, in the */tmp* directory. You can specify the location of the log files.

The `db2setup.log` file captures all DB2 installation information including errors. The `db2setup.err` file captures any error output that is returned by Java (for example, exceptions and trap information).

There is no longer a `db2setup.his` file. Instead, the DB2 installer saves a copy of the `db2setup.log` file in the `DB2_DIR/install/logs/` directory, and renames it to `db2install.history`. If `db2install.history` already exists, the copy of the `db2setup.log` will be renamed `db2install.history.xxxx` where `xxxx` are digits starting at 0000 and increasing by 1 for the next log file.

Note: You may need to run the `eject` command or unmount the CD drive before you can physically remove the installation media.

Chapter 42. Installing the DB2 Information Center using the DB2 Setup wizard (Windows)

Using the DB2 Setup wizard, you can define your installation preferences and install the DB2 Information Center on a computer that uses the Windows operating system.

Before you begin

This section lists the hardware, operating system, software, and communication requirements for installing the *DB2 Information Center* on Windows.

Operating system requirements

You require one of the following operating systems:

- Windows Vista
- Windows XP
- Windows Server 2003
- Windows Server 2008

The *DB2 Information Center* will run on Windows and Linux on AMD/EMT 64, but it does not exploit the 64-bit architecture.

Software requirements

You require one of the following browsers:

- Firefox 1.0 and higher
- Internet Explorer (IE) 6.0 and higher
- Mozilla-based browsers 1.7 and higher
- Safari 3.0
- Konqueror (UI base mode only). Base mode restricts the users basic functionality such as displaying topics, locating topics in the table of contents, and searching for topics.

Communication requirements

You require TCP/IP protocol.

About this task

DB2 database product documentation can be accessed in three ways:

- at the IBM Web site
- on an intranet server
- on a version installed on your computer

By default, DB2 database products access DB2 documentation at the IBM Web site. If you want to access the DB2 documentation on an intranet server or on your own computer, you must install the DB2 documentation from the *DB2 Information Center DVD*, or from an image downloaded from Passport Advantage, or DB2 database product documentation website.

Restrictions

- To install a network version of the *DB2 Information Center*, you require an account with administrative privileges to install. To install a Workstation version of the DB2 Information Center, you do not need administrative privileges.

- You cannot install the *DB2 Information Center* in a location where a DB2 database product is installed. Similarly, the *DB2 Information Center* cannot coexist in the same installation path of a previous version of the *DB2 Information Center* on the same system.
- You can only install one copy of each version of the *DB2 Information Center* on your system. For example, you can install a copy of the *DB2 Information Center Version 9.7* and *DB2 Information Center Version 10.1* on the same system, but you cannot install two copies of the *DB2 Information Center Version 10.1*.
- If you install the *DB2 Information Center* on a system with a firewall, and you plan for other systems to access the *DB2 Information Center*, you must open the port in your firewall setting.

Procedure

To install the *DB2 Information Center* using the DB2 Setup wizard:

1. Log on to the system with the account that you have defined for the *DB2 Information Center* installation.
2. If you have the *DB2 Information Center product DVD*, insert your DB2 database product DVD into the drive. If enabled, the auto-run feature automatically starts the DB2 Setup launchpad. If the auto-run does not work, use Windows Explorer to browse the DB2 database product DVD.
3. If you downloaded the image from an online source (such as Passport Advantage, or, ibm.com), uncompress the *DB2 Information Center* installation image.
4. Double-click the **setup** icon.
5. From the DB2 Setup Launchpad, you can view installation prerequisites and the release notes, or you can proceed directly to the installation. You should review the installation prerequisites and release notes for late-breaking information.
6. Click **Install a Product** and the Install a Product window displays.
7. On the Install a Product window, if you do not have an existing *DB2 Information Center* installed on your computer, launch the installation by clicking **Install New**.
8. On the Welcome to the DB2 Setup wizard window, click **Next**. The DB2 Setup wizard will guide you through the program setup process. The DB2 Setup wizard will determine the system language, and launch the setup program for that language. Online help is available to guide you through the remaining steps. To invoke the online help, click **Help** or press F1. You can click **Cancel** at any time to end the installation.

Results

Your *DB2 Information Center* will be installed, by default, in the *Program_Files\IBM\DB2 Information Center\Version 10.1* directory, where *Program_Files* represents the location of the Program Files directory.

For information about errors encountered during installation, review the installation log file located, by default, in the *My Documents\DB2LOG* directory. You can specify the location of the log files. The log file uses the following format: *DB2-DOCE-DateTime.log*, for example, *DB2-DOCE-Wed Apr 11 08_35_2007.log*

Chapter 43. DB2 Information Center installation options

The DB2 Information Center can be installed locally so that you can access it when you are not connected to the internet.

The DB2 Information Center can be accessed from:

- the ibm.com website
- a server on your internal network
- a copy installed on your computer

By default, DB2 products access the DB2 Information Center at the ibm.com website: <http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/>. However, if you want to access the DB2 Information Center on an intranet server or on your own computer, you must install the DB2 Information Center.

The following types of DB2 Information Centers are available:

- The regular DB2 Information Center
This is the same DB2 Information Center that is distributed with DB2 database products. It comes with an installer and other programs that let you install the Information Center on your computer. The installation program requires that you have administrative authority on your computer to complete the installation.
- A Workstation version of the DB2 Information Center
These packages allow you to run the DB2 Information Center on your computer if you do not have administrator or root authority. The Workstation version of the DB2 Information Center run in "stand-alone" mode; it cannot be accessed by other clients on your network. There are no services or daemons associated with this type of DB2 Information Center, therefore you must start and stop it manually. It also differs from the regular DB2 Information Center because it determines the locale from the system locale of the system, not from the browser.

You can install either of the DB2 Information Centers with the DB2 Information Center DVD found in your product Media Pack. Alternatively, you can download the DB2 Information Center installation image from IBM Passport Advantage: <http://www.ibm.com/software/passportadvantage>.

The following table provides recommendations on possible options for accessing DB2 product documentation in the DB2 Information Center based on your work environment.

Internet access	Intranet access	Recommendation
Yes	Yes	Access the DB2 Information Center at the ibm.com website, or access the DB2 Information Center installed on an intranet server.
Yes	No	Access the DB2 Information Center at the ibm.com website.
No	Yes	Access the DB2 Information Center installed on an intranet server.

Internet access	Intranet access	Recommendation
No	No	Access the DB2 Information Center on a local computer, or access the stand-alone Workstation version of the DB2 Information Center..

Accessing the DB2 Information Center at the [ibm.com](http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/index.jsp) website

The DB2 Information Center can be found on the Internet at <http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/index.jsp>. If you configure your local DB2 installation to use this version of the DB2 Information Center, you will always be accessing the most up-to-date information about the DB2 product.

On Windows platforms, use the **db2set** command to configure the registry variables on your computer to access the DB2 Information Center from either the [ibm.com](http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/index.jsp) website or from your computer. You can also change these variables through the DB2 GUI tools, if you have installed the DB2 GUI tools on your system.

Accessing the DB2 Information Center on your computer

After installing your DB2 product, you need to install the DB2 Information Center if you want to access DB2 documentation on this computer. Use the DB2 Information Center DVD in the media pack or the files that you downloaded from IBM Passport Advantage to install the DB2 Information Center.

Accessing the DB2 Information Center on an intranet server

You can install a copy of the DB2 Information Center on an intranet server so that everyone on the intranet can have access to DB2 documentation without having to install the documentation on every machine.

Using a response file for product installation, you can configure every IBM data server client to access the documentation on the intranet server. The response file must contain the configuration information necessary to access the DB2 Information Center with the host name and the port number of the DB2 Information Center on the intranet server. This configuration can also be done by using the **db2setup** command and choosing a custom installation from any of the DB2 server or IBM data server client products.

You can also change the settings for already-installed IBM Data Server Client to use the DB2 Information Center hosted on your Intranet. To change the DB2 Information Center registry variables on each computer, use the **db2set** command.

The **DB2_DOCHOST** registry variable contains the host name, and **DB2_DOCPORT** is the registry variable used for the port number. These values must match the host and port on the intranet server where the DB2 Information Center is installed.

Chapter 44. Starting or stopping the Information Center (Linux and Windows)

The DB2 Information Center is started and stopped using either the Windows services dialog, or a Linux daemon.

About this task

On Linux operating systems, the *DB2 Information Center* daemon is part of the *DB2 Information Center* installation. This daemon is a background process that runs the Information Center. The daemon is initialized by the script **db2icd** which can be found in *INST_PATH/doc/bin* where *INST_PATH* is the installation path for your DB2 database product. The only time you should need to start or stop the daemon manually is when you want to change the configuration variables for the daemon. Normally, the daemon is started at system startup, according to the run levels created during the installation of the *DB2 Information Center*. The Workstation version of the DB2 Information Center does not have a daemon.

On Windows operating systems, to stop or start the Information Center use the services control panel applet: **Control Panel > Administrative Tools > Services**. Right-click the service labelled *DB2 Information Center*. Select either **Stop** or **Start**.

Procedure

On Linux operating systems, to stop and start the Information Center daemon:

1. Log in as a user with root user authority.
2. Halt the daemon if it is already running. At a command line, enter:

```
INIT_DIR/db2icd stop
```

where *INIT_DIR* is the installation directory of the *db2icd* file listed previously.

3. Change any of the variables for the daemon by editing the *db2ic.conf* file. Currently, you can modify the TCP port number for the documentation service, and the location of the temporary workspace used by the daemon while it is running.
4. Start the daemon. At a command line, enter:

```
INIT_DIR/db2icd start
```

where *INIT_DIR* is the installation directory of the *db2icd* file listed previously.

Results

When the daemon restarts, it uses the new environment variables.

There is also an option to shut down and restart the daemon immediately. At a command line, enter:

```
INIT_DIR/db2icd restart
```

where *INIT_DIR* is the installation directory of the *db2icd* file listed previously.

You can check the status of the daemon at any time. At a command line, enter:

```
INIT_DIR/db2icd status
```

where *INIT_DIR* is the installation directory of the db2icd file listed previously. The daemon returns the current status, and displays the process ID of the daemon if it is active.

Chapter 45. Manually updating the DB2 Information Center installed on your computer or intranet server

If you have installed the DB2 Information Center locally, you can obtain and install documentation updates from IBM.

About this task

Updating your locally installed *DB2 Information Center* manually requires that you:

1. Stop the *DB2 Information Center* on your computer, and restart the Information Center in stand-alone mode. Running the Information Center in stand-alone mode prevents other users on your network from accessing the Information Center, and allows you to apply updates. The Workstation version of the DB2 Information Center always runs in stand-alone mode. .
2. Use the Update feature to see what updates are available. If there are updates that you must install, you can use the Update feature to obtain and install them

Note: If your environment requires installing the *DB2 Information Center* updates on a machine that is not connected to the internet, mirror the update site to a local file system by using a machine that is connected to the internet and has the *DB2 Information Center* installed. If many users on your network will be installing the documentation updates, you can reduce the time required for individuals to perform the updates by also mirroring the update site locally and creating a proxy for the update site.

If update packages are available, use the Update feature to get the packages. However, the Update feature is only available in stand-alone mode.

3. Stop the stand-alone Information Center, and restart the *DB2 Information Center* on your computer.

Note: On Windows 2008, Windows Vista (and higher), the commands listed later in this section must be run as an administrator. To open a command prompt or graphical tool with full administrator privileges, right-click the shortcut and then select **Run as administrator**.

Procedure

To update the *DB2 Information Center* installed on your computer or intranet server:

1. Stop the *DB2 Information Center*.
 - On Windows, click **Start > Control Panel > Administrative Tools > Services**. Then right-click **DB2 Information Center** service and select **Stop**.
 - On Linux, enter the following command:
`/etc/init.d/db2icdv10 stop`
2. Start the Information Center in stand-alone mode.
 - On Windows:
 - a. Open a command window.
 - b. Navigate to the path where the Information Center is installed. By default, the *DB2 Information Center* is installed in the `Program_Files\IBM\DB2 Information Center\Version 10.1` directory, where *Program_Files* represents the location of the Program Files directory.

- c. Navigate from the installation directory to the doc\bin directory.
- d. Run the help_start.bat file:
`help_start.bat`
- On Linux:
 - a. Navigate to the path where the Information Center is installed. By default, the *DB2 Information Center* is installed in the /opt/ibm/db2ic/V10.1 directory.
 - b. Navigate from the installation directory to the doc/bin directory.
 - c. Run the help_start script:
`help_start`

The system's default Web browser opens to display the stand-alone Information Center.

3. Click the **Update** button (🔄). (JavaScript must be enabled in your browser.) On the right panel of the Information Center, click **Find Updates**. A list of updates for existing documentation displays.
4. To initiate the installation process, check that the selections you want to install, then click **Install Updates**.
5. After the installation process has completed, click **Finish**.
6. Stop the stand-alone Information Center:
 - On Windows, navigate to the doc\bin directory within the installation directory, and run the help_end.bat file:
`help_end.bat`
 - Note:** The help_end batch file contains the commands required to safely stop the processes that were started with the help_start batch file. Do not use Ctrl-C or any other method to stop help_start.bat.
 - On Linux, navigate to the doc/bin directory within the installation directory, and run the help_end script:
`help_end`
 - Note:** The help_end script contains the commands required to safely stop the processes that were started with the help_start script. Do not use any other method to stop the help_start script.
7. Restart the *DB2 Information Center*.
 - On Windows, click **Start > Control Panel > Administrative Tools > Services**. Then right-click **DB2 Information Center** service and select **Start**.
 - On Linux, enter the following command:
`/etc/init.d/db2icdv10 start`

Results

The updated *DB2 Information Center* displays the new and updated topics.

Part 7. Converting instances to a new DB2 database product

You can convert from one DB2 database product to another DB2 database product or convert from one DB2 copy to another DB2 copy.

Chapter 46. Converting instances from one DB2 database product to another (Windows)

To convert from one DB2 database product to another (or from one DB2 copy to another) on Windows operating systems, use the **Work with existing** option in the DB2 database product setup wizard. Once the installation is complete, you can update the licensing information and uninstall the previous product.

Before you begin

Your system must comply with the system prerequisites for the new DB2 database product that you are installing.

If there were any fix packs installed on the original DB2 database product, you should obtain the image for that fix pack level and install that directly, as opposed to installing the DB2 database product and then applying the fix pack. The reason for this is that each DB2 fix pack image is a full refresh image.

About this task

These instructions apply only when the old and new DB2 database products share the same version level (for example, Version 10.1).

Procedure

To convert between DB2 database products:

1. Log on as a user with administrator access.
2. Stop all DB2 processes and services.
3. Insert the DB2 database product DVD for the new edition you are installing into the drive. The IBM DB2 database product launchpad will open.
4. The DB2 database product setup wizard will determine the system language, and will start the setup program for that language. To start the DB2 database product setup wizard manually:
 - a. Click **Start** and select the **Run** option.
 - b. In the **Open** field, type the following command:

```
x:\setup /i language
```

where *x* represents your DVD drive and *language* is the territory identifier for your language, for example, EN for English.

If the */i* flag is not specified, the installation program will run in the default language of the operating system.

Click **OK**.

5. Choose **Work with existing** to start the installation and proceed by following the setup program's prompts. Online help is available to guide you through these steps. To invoke the online help, click **Help** or press F1. You can click **Cancel** at any time to end the installation.
6. Add the license file for the new DB2 database product by issuing the **db2licm** command.

```
db2licm -a LicenseFile
```

The license file is found in the `db2\license` directory on the product activation CD.

7. Remove the original DB2 database product's license by issuing the **db2licm** command.
 - a. Issue the command **db2licm -l** to find the original DB2 database product's *product-identifier* value.
 - b. Issue the following command to remove the license:
`db2licm -r product-identifier`
 - c. Issue the command **db2licm -l** and check the expiry date value. If the original DB2 database product does not have an expiry date of Expired, repeat the **db2licm -r** command until all of the original DB2 database product's licenses are removed.
8. Remove the original DB2 database product using the Add or Remove Control Programs Control Panel window.
9. When the original DB2 database product is completely uninstalled, reactivate the administration server and the services using the **db2start** command or the Services window of the Control Panel.

Results

When these steps are completed, your instances are configured to run in the new DB2 database product installation.

Chapter 47. Converting instances from one DB2 database product to another (Linux and UNIX)

To convert from one DB2 database product to another (or from one DB2 copy to another) on Linux or UNIX operating systems, install the product, update the licensing information, then reconfigure the instances.

Before you begin

Your system must comply with the system prerequisites for the DB2 database product edition that you are installing.

About this task

Be careful when you issue the **db2_deinstall** command. If you issue the command without specifying appropriate command options, you might remove all DB2 database product components.

These instructions apply only when switching among the same version of a DB2 database product.

Procedure

To convert between DB2 database products:

1. Complete the DB2 database product installation for the new edition.
You do not need to create an instance. Existing instances can be re-configured by issuing **db2iupdt** or **db2nrupdt** commands when the installation is complete.
2. Add the license file for the new database product by issuing the **db2licm** command.

```
db2licm -a LicenseFile
```

The license file is found in the `db2/license` directory on the product activation CD.

3. To obtain a list of the names of all instances on your system, run the **db2ilist** command:

```
DB2DIR/db2ilist
```

Where *DB2DIR* represents the location where the original DB2 database product is installed. By default, it is `/opt/ibm/db2/V10.1` on Linux and `/opt/IBM/db2/V10.1` on UNIX operating systems.

4. Stop any of the instances that are running.
5. To re-configure the instances, issue the following commands:
 - **db2iupdt** for each root instance
 - **db2nrupdt** for the non-root instance.

For example:

```
DB2DIR/bin/db2iupdt InstanceName
```

where *DB2DIR* represents the location where the new DB2 database product is installed and *InstanceName* represents the name of the instance.

6. Remove the original DB2 database product license by running the **db2licm** command.
 - a. Run the command **db2licm -l** to find the original DB2 database product *product-identifier* value.
 - b. Issue the following command to remove the license:
db2licm -r product-identifier
 - c. Issue the command **db2licm -l** and check the expiry date value. If the original DB2 database product does not have an expiry date of Expired, repeat the **db2licm -r** command until you have successfully removed all of the original licenses for the DB2 database products.
7. Remove the product signature of the original DB2 database product:
 - a. Find the product signature by issuing the **db2ls** command:
db2ls -q -a -b DB2DIR

Where *DB2DIR* represents the location where the new DB2 database product is installed.
You will see product signatures for both the old and new DB2 database products. For example, EXPRESS_PRODUCT_SIGNATURE or WSE_PRODUCT_SIGNATURE.
 - b. Remove the old product signature by issuing the **db2_deinstall** command:
DB2DIR/install/db2_deinstall -F product_signature

Where *DB2DIR* represents the location where the new DB2 database product is installed. You will no longer see the old product signature in the output from the **db2ls** command.
8. Optional: Uninstall the original DB2 database product.

Results

When these steps are completed, your instances are configured to run in the new DB2 database product installation.

Part 8. Convert instances to the IBM DB2 pureScale Feature

Converting to the IBM DB2 pureScale Feature might require that you configure some of your DB2 environment components if you want them to run in the new environment.

Your DB2 environment has several components, such as DB2 servers, DB2 clients, database applications, and routines. To successfully convert your environment to the IBM DB2 pureScale Feature, you must perform a variety of tasks.

Chapter 48. Planning to convert your environment to DB2 pureScale

Planning to convert your environment to DB2 pureScale requires that you review all the applicable conversion prerequisites, pre-conversion tasks, conversion tasks, and post-conversion tasks.

Procedure

1. Review the Installation prerequisites for the DB2 pureScale Feature
2. Review all prerequisites for the task, especially obtaining root access and required DB2 authorization. Refer to the documentation on the specific commands for the authorization required.
3. Review the Chapter 49, "Pre-conversion tasks for DB2 pureScale environments," on page 443 topic to prepare for enabling your DB2 pureScale environment. Pre-conversion tasks are:
 - Verifying that your databases are ready for conversion
 - Converting table spaces
 - Setting up GPFS
 - Backing up your databases and configuration information
 - Taking the DB2 server offline
4. Convert your DB2 pureScale environment as described in Chapter 50, "Converting your existing DB2 instances to a DB2 pureScale environment," on page 449.
5. After converting to a DB2 pureScale environment, you must perform several post-conversion tasks to help ensure that your DB2 servers perform at their optimum level. Review the Chapter 51, "Post-conversion tasks for a DB2 pureScale environment," on page 451 topic which describes tasks such as
 - Adjust database configuration parameters
 - Activate your database
 - Verify the conversion
 - Backing up your databases and configuration information

Chapter 49. Pre-conversion tasks for DB2 pureScale environments

Before you convert your environment to DB2 pureScale, review the restrictions to identify anything that can affect the conversion. You must address any issues before converting to have a successful conversion.

Procedure

Prepare for converting your environment to DB2 pureScale by performing the following tasks:

1. Set up a GPFS file system. See “Setting up a GPFS file system for a DB2 pureScale environment.” You must ensure that all data and logs are on a DB2 cluster file system before you convert.
2. Convert table spaces. Refer to “Converting table spaces for a DB2 pureScale environment” on page 445. You must ensure that you have automatic table spaces before you convert.
3. Verify that your databases are ready for converting to a DB2 pureScale environment. Refer to “Verifying that your databases are ready for a DB2 pureScale environment” on page 446. Verifying that your databases are ready for conversion identifies any problems before the actual conversion. You must resolve any problems found before you proceed with the conversion.
4. Back up your databases to be able to restore them in the original pre-converted system if needed. For details, see “Backing up databases when converting to a DB2 data sharing environment” on page 447.
5. Back up DB2 server configuration and diagnostic information. Backing up the configuration and diagnostic information gives you a record of your current configuration that you can compare with the configuration after the conversion. You can also use this information to create new instances or databases using the same configuration that you had before conversion. Refer to “Backing up DB2 server configuration and diagnostic information” in *Upgrading to DB2 Version 10.1*.
6. If the error capture level of the diagnostic log files is set to 2 or less, set the error capture level to level 3 or higher. In this release, all significant conversion events are logged in the diagnostic log when the **diaglevel** database manager configuration parameter is set to 3 (default value) or higher.
For details, see “Setting the diagnostic log file error capture level” in *Troubleshooting and Tuning Database Performance*.
7. Take the DB2 server offline before conversion. Refer to “Taking a DB2 server offline for upgrade or for conversion to a DB2 pureScale environment” in *Upgrading to DB2 Version 10.1*.

Setting up a GPFS file system for a DB2 pureScale environment

To set up an IBM General Parallel File System (GPFS) file system for a DB2 pureScale environment, you can create a new GPFS file system, GPFS file system, or have DB2 cluster services take control of an existing GPFS file system.

Before you begin

- You must have installed the IBM DB2 pureScale Feature.
- Ensure that you have root access to use the **db2cluster_prepare** command.

About this task

In a DB2 pureScale environment, all data and logs must be on an IBM General Parallel File System (GPFS). If **db2checkSD** reported any table spaces that are not managed by automatic storage, or that are not on GPFS, you must convert the table spaces and move the data to GPFS before conversion. Before moving the data, you might need to create a DB2 managed GPFS cluster and file system, or verify that an existing GPFS cluster and file system can be used by DB2 as a user managed GPFS.

Procedure

To configure a GPFS file system for a DB2 pureScale environment:

Choose one of the following options:

- Set up a new DB2 cluster file system on a host that does not have an existing GPFS cluster. This is the recommended method to set up a DB2 pureScale environment.

1. Use the **db2cluster_prepare** command to create a DB2 managed GPFS file system. For example:

```
db2cluster_prepare -instance_shared_dev /dev/hdisk1
```

2. Verify the mount location. To list the file systems, run the following **db2cluster** command. For example:

```
db2cluster -cfs -list -filesystem
```

- Use an existing user managed GPFS cluster and have DB2 cluster services take over the management of the cluster. Run the **db2cluster_prepare** command with the **-cfs_takeover** option to have DB2 cluster services manage an existing GPFS cluster.

```
db2cluster_prepare -cfs_takeover
```

- Identify and continue using an existing user managed GPFS cluster.

Perform these steps to use an existing GPFS as a user managed file system. DB2 pureScale Feature can use a user managed GPFS cluster and file system, but when the DB2 installer and instance utilities are used to extend a DB2 pureScale instance that has a user managed GPFS cluster and file system, the host must be added to the GPFS cluster manually and the related file systems on the host must be mounted manually. Similarly, when the DB2 installer and instance utilities are used to shrink a DB2 pureScale instance that has a user managed GPFS cluster and file system, the host must be removed from the GPFS cluster manually and the related file systems on the host must be unmounted manually. When the GPFS cluster and file system are DB2 managed, the DB2 installer and instance utilities will add or remove the host and issue the required mount or unmount commands when the host is added to or removed from the DB2 pureScale instance. Before using an existing GPFS you must verify that GPFS is set up correctly:

1. Verify the cluster. For example, run the following from the <DB2InstallPath>/bin directory:

```
db2cluster -cfs -verify -configuration
```

2. Verify the file system. You must perform this step for each file system that you plan to use with DB2 pureScale Feature. For example, run the **db2cluster** command from the <DB2InstallPath>/bin directory with the following options:

```
db2cluster -cfs -verify -configuration
```

Converting table spaces for a DB2 pureScale environment

Converting your environment to DB2 pureScale requires changes to table spaces.

Before you begin

- You must have installed the DB2 pureScale Feature.

About this task

In a DB2 pureScale environment, table spaces must be managed by automatic storage. If **db2checkSD** reported any table spaces that are not managed by automatic storage, or that are not on GPFS, you must convert the table spaces and move the data to GPFS before converting. The process steps you take differ, depending on whether you want to have a DB2 managed GPFS or have a user managed GPFS.

Procedure

Convert tables spaces and move data. The steps required to convert to automatic storage table spaces vary depending on how your existing table spaces are defined:

Table 53. Convert tables spaces

Scenario	Actions
Case 1: Catalog table space is system managed space (SMS)	Recreate the database on the GPFS created in the previous step and reload it. Use the COPY option of the db2move command to copy one or more schemas from the original database to the new database.

Table 53. Convert tables spaces (continued)

Scenario	Actions
Case 2: <ul style="list-style-type: none"> • Catalog table space is not SMS • May have SMS user table spaces • May have some SMS temporary table spaces • Not all database managed (DMS) table spaces are managed by automatic storage 	<ul style="list-style-type: none"> • If the database is not enabled for automatic storage, use the ALTER DATABASE command with the ADD STORAGE ON storage-path option to enable it for automatic storage • If SMS table spaces exist: <ul style="list-style-type: none"> – drop them – re-create them as table spaces that are managed by automatic storage – reload them • If SMS user temporary table spaces exist: <ul style="list-style-type: none"> – drop them – re-create them as table spaces that are managed by automatic storage – re-create any global temporary tables in the new table space as dropping a user temporary table space will also drop any global temporary tables in that table space • Perform a redirected restore to create a new copy of the database on the GPFS storage path. During the redirected restore, specify SET TABLESPACE CONTAINERS FOR tablespace-id USING AUTOMATIC STORAGE for any DMS table spaces, to convert them to automatic storage.
Case 3: All table spaces are managed by automatic storage but are not on GPFS	Perform a redirected restore to create a new copy of the database on the GPFS storage path.

Verifying that your databases are ready for a DB2 pureScale environment

Before converting your environment to DB2 pureScale, use the **db2checkSD** command to verify that your databases are ready for use in a DB2 pureScale environment.

Before you begin

- Review the topics on the **db2checkSD** commands and ensure that you have the authorization required to run them.
- Ensure that all the local databases that you want to convert are cataloged.
- Ensure that you meet the installation requirements for the DB2 pureScale Feature.
- Install a DB2 Version 10.1 copy so you can run the **db2checkSD** command.

Procedure

To verify that your databases are ready for conversion:

1. Log on as the owner of the DB2 instance that you want to convert.

2. If the instance owning the databases that you want to verify is not running, start the instance by running the **db2start** command.
3. From a DB2 command-line prompt change to the `$DB2DIR/bin` directory where `DB2DIR` is the location that you specified during the DB2 installation:

```
cd $DB2DIR/bin
```
4. Check that your databases meet the requirements for converting to a DB2 pureScale environment by running the **db2checkSD** command. The following example shows how to run the **db2checkSD** command on the SAMPLE database:

```
db2checkSD SAMPLE -l db2checkSD.log -u adminuser -p password
DBT5000I The db2checkSD command completed successfully.
The specified database can be upgraded to a DB2 pureScale environment.
```

where `db2checkSD.log` is the log file created in the `INSTHOME/sql11ib/db2dump/` directory and `INSTHOME` is the instance home directory. The `db2checkSD.log` includes details on errors and warnings. You must correct these errors before you convert.

5. Ensure that the log file for the **db2checkSD** command starts with the following text to confirm that you are running the correct command level:

```
Version of DB2CHECKSD being run: Version 10.1.
```

Backing up databases when converting to a DB2 data sharing environment

Before and after the process of converting your environment to DB2 pureScale, it is strongly recommended that you perform a full *offline* database backup.

Before you begin

- To backup a database, you require SYSADM, SYSCTRL, or SYSMAINT authority.
- Databases must be cataloged. To view a list of all the cataloged databases in the current instance, enter the following command:

```
db2 LIST DATABASE DIRECTORY
```

Procedure

To perform a full offline back up for each of your local databases:

1. Disconnect all applications and users from the database. To get a list of all database connections for the current instance, issue the **LIST APPLICATIONS** command:

```
db2 LIST APPLICATIONS
```

If all applications are disconnected, this command returns the following message:

```
SQL1611W No data was returned by the Database System Monitor.
SQLSTATE=00000
```

To disconnect all applications and users, use the **FORCE APPLICATION** command:

```
db2 FORCE APPLICATION ALL
```

2. Backup your database using the **BACKUP DATABASE** command. The following is an example for UNIX operating systems:

```
db2 BACKUP DATABASE database_alias USER username USING password TO backup-dir
```

where *database_alias* is the database alias, the user name is *username*, the password is *password*, and the directory to create back up files is *backup-dir*.

In partitioned database environments, back up all database partitions. For details, see “Backing up partitioned databases” in *Data Recovery and High Availability Guide and Reference*.

If you activated and configured DB2 Advanced Copy Services (ACS) on your databases in DB2 Version 9.5 or later, you can use the **USE SNAPSHOT** parameter to perform a snapshot backup. However, you can only restore a snapshot backup to an instance of the same version. You cannot use snapshot backup to move to a new server. For details, see .Performing a snapshot backup in *Data Recovery and High Availability Guide and Reference*

If you performed a full offline database backup recently and you cannot perform another one before conversion, you can perform an incremental offline database backup instead.

3. Optional: Test the integrity of a backup image to ensure that the image can be restored using the **db2ckbkp** command. The following command is an example on UNIX operating systems:

```
cd backup-dir
db2ckbkp SAMPLE.0.arada.NODE0000.CATN0000.20091014114322.001
```

```
[1] Buffers processed: #####
```

```
Image Verification Complete - successful.
```

Chapter 50. Converting your existing DB2 instances to a DB2 pureScale environment

Converting your existing DB2 instances to DB2 pureScale requires that you update your instance to a DB2 pureScale instance.

Before you begin

- Review the information in Chapter 48, “Planning to convert your environment to DB2 pureScale,” on page 441.
- Carry out the steps in Chapter 49, “Pre-conversion tasks for DB2 pureScale environments,” on page 443.
- If your database is using the DB2 Partitioned Database Environment, you must convert your database to a single partition database before enabling to DB2 pureScale.

About this task

This task describes the steps for converting your existing instances to a DB2 pureScale environment. Some steps in this topic require that you are logged in as root.

Procedure

1. If you have not already done so, install the DB2 pureScale Feature.
2. As root, convert the instance to a DB2 pureScale instance by issuing the **db2iupdt** command.
3. Verify the instance conversion is successful by connecting to the databases and issuing a small query, as shown in the following example:

```
db2 connect to sample
```

```
Database Connection Information
```

```
Database server          = DB2/AIX64 10.1.0  
SQL authorization ID     = TESTDB2  
Local database alias     = SAMPLE
```

```
db2 select * from syscat.dbauth
```

4. Add additional members or cluster caching facilities by using the **db2iupdt** command or **db2isetaup** command. It is recommended that you have two cluster caching facilities.

What to do next

Perform the steps in the Chapter 51, “Post-conversion tasks for a DB2 pureScale environment,” on page 451 topic.

Chapter 51. Post-conversion tasks for a DB2 pureScale environment

After converting your environment to DB2 pureScale, you must perform several post-conversion tasks to help ensure that your DB2 servers perform at their optimum level.

Procedure

To help ensure that your DB2 servers perform at their optimum level:

1. If you set the **diaglevel** database manager configuration parameter to 3 or higher as part of the pre-conversion tasks for a DB2 pureScale environment, reset this parameter to the value that was set before the conversion.
2. Adjust your database configuration settings to meet new requirements in a DB2 pureScale environment. See “Adjusting database configuration parameters to meet DB2 pureScale environment requirements” for steps that you must take to meet new requirements in a DB2 pureScale environment.
3. Ensure that the libraries for your external routines remain in the location where they were before the conversion. If necessary, restore these libraries from the backups that you performed. Refer to “Backing up DB2 server configuration and diagnostic information” in *Upgrading to DB2 Version 10.1*.
4. Start your database and all necessary database services. For details, see “Activating a database after upgrade” in *Upgrading to DB2 Version 10.1*.
5. If you created local or global file or pipe event monitors in Version 9.7 and the directory or named pipe that you specified is not on a shared file system, drop and re-create them indicating that the directory or named pipe is on a shared file system. Otherwise, these event monitors return an error during member failover. For details, see the CREATE EVENT MONITOR statement documentation.
6. To use the new high availability function in your global event monitors and the default DB2DETAILDEADLOCK event monitor, drop them and re-create them. For details, see the CREATE EVENT MONITOR statement documentation.
7. Back up your databases.
8. Verify that your environment conversion is successful by connecting to the databases and issuing a small query, as shown in the following example:

```
db2 connect to sample

      Database Connection Information

      Database server      = DB2/AIX64 10.1.0
      SQL authorization ID = TESTDB2
      Local database alias = SAMPLE

db2 select * from syscat.dbauth
```

Adjusting database configuration parameters to meet DB2 pureScale environment requirements

After converting your environment for DB2 pureScale, you must make a number of changes to database configuration parameters.

Procedure

1. Ensure that there is sufficient disk space to hold the logs retrieved from members during a log merge. Providing space for two log files per member in the log path, mirror log path, or overflow log path allows one log file to be kept while another log file is being retrieved. Having space for two files helps to avoid the problem of a log merge having to wait for log files to be retrieved.
2. To support the additional memory requirements for group restart, rollforward, log merges, and log file retrieval, ensure that the **dbheap** database configuration parameter is set to AUTOMATIC. You should also use the AUTOMATIC setting if you are making use of table partitioning; this takes into account the memory used to synchronize table partitioning data between the cluster caching facility (CF) and all members.
3. To support the additional locking requirements for data coherency between members, ensure that the **locklist** parameter is set using the XX AUTOMATIC option, where XX is 3% of the total number of pages for all the existing bufferpools.
4. If you have large buffer pools, you have not enabled self-tuning memory, and your DB2 server is using nearly all the physical memory on the system, consider decreasing the buffer pool size or enabling self-tuning memory.
5. Review the memory considerations for restart light and make any necessary adjustments.

Backing up databases when converting to a DB2 data sharing environment

Before and after the process of converting your environment to DB2 pureScale, it is strongly recommended that you perform a full *offline* database backup.

Before you begin

- To backup a database, you require SYSADM, SYSCTRL, or SYSMAINT authority.
- Databases must be cataloged. To view a list of all the cataloged databases in the current instance, enter the following command:

```
db2 LIST DATABASE DIRECTORY
```

Procedure

To perform a full offline back up for each of your local databases:

1. Disconnect all applications and users from the database. To get a list of all database connections for the current instance, issue the **LIST APPLICATIONS** command:

```
db2 LIST APPLICATIONS
```

If all applications are disconnected, this command returns the following message:

```
SQL1611W No data was returned by the Database System Monitor.  
SQLSTATE=00000
```

To disconnect all applications and users, use the **FORCE APPLICATION** command:

```
db2 FORCE APPLICATION ALL
```

2. Backup your database using the **BACKUP DATABASE** command. The following is an example for UNIX operating systems:

```
db2 BACKUP DATABASE database_alias USER username USING password TO backup-dir
```

where *database_alias* is the database alias, the user name is *username*, the password is *password*, and the directory to create back up files is *backup-dir*.

In partitioned database environments, back up all database partitions. For details, see “Backing up partitioned databases” in *Data Recovery and High Availability Guide and Reference*.

If you activated and configured DB2 Advanced Copy Services (ACS) on your databases in DB2 Version 9.5 or later, you can use the **USE SNAPSHOT** parameter to perform a snapshot backup. However, you can only restore a snapshot backup to an instance of the same version. You cannot use snapshot backup to move to a new server. For details, see .Performing a snapshot backup in *Data Recovery and High Availability Guide and Reference*

If you performed a full offline database backup recently and you cannot perform another one before conversion, you can perform an incremental offline database backup instead.

3. Optional: Test the integrity of a backup image to ensure that the image can be restored using the **db2ckbkp** command. The following command is an example on UNIX operating systems:

```
cd backup-dir
db2ckbkp SAMPLE.0.arada.NODE0000.CATN0000.20091014114322.001
```

```
[1] Buffers processed: #####
```

```
Image Verification Complete - successful.
```

Part 9. Applying fix packs

It is recommended that you keep your DB2 database environment running at the latest fix pack level to ensure problem-free operation. To install a fix pack successfully, perform all of the necessary preinstallation and post-installation tasks.

About this task

A DB2 fix pack contains updates and fixes for problems (Authorized Program Analysis Reports, or "APARs") found during testing at IBM, as well as fixes for problems reported by customers. The APARLIST.TXT file describes the fixes contained in each fix pack and it is available for download at <ftp://ftp.software.ibm.com/ps/products/db2/fixes/english-us/aparlist/>.

Fix packs are cumulative. This means that the latest fix pack for any given version of DB2 database contains all of the updates from previous fix packs for the same version of DB2 database.

The fix pack images available are:

- A single server image.

The single server image contains the new and updated code required for all DB2 database server products and the IBM Data Server Client. If more than one DB2 database server product is installed in a single location, the DB2 database server fix pack applies maintenance code updates to all the installed DB2 database server products. The Data Server Client fix pack is contained within the one DB2 database server fix pack (namely the fix pack that can service any one of the following database server products: DB2 Enterprise Server Edition, DB2 Workgroup Server Edition, DB2 Express Edition, DB2 Connect Enterprise Edition, DB2 Connect Application Server Edition, DB2 Connect Unlimited Edition for zSeries, and DB2 Connect Unlimited Edition for i5/OS). You can use the DB2 database server fix pack to upgrade a Data Server Client.

A single server image can also be used to install any of the DB2 database server products, at a particular fix pack level, with a DB2 try and buy license by default.

The single server fix pack image contains DB2 try-and-buy licenses for all DB2 server products. When you select a new DB2 server product to install or a previously installed DB2 server product to update, the try-and-buy licenses are installed. The try-and-buy licenses do not affect any valid licenses already installed in the same DB2 installation path. Regarding DB2 Connect server products, if you run the **db21icm -l** command to query valid licenses, the try-and-buy license for DB2 Connect server product might display as an invalid license. However, if you do not need to use the DB2 Connect functionality, you can ignore the report. To remove the try-and-buy license for DB2 Connect server, use the **db21icm** command.

- A fix pack for each of the other DB2 database products.

Use this fix pack only if you only have non-server database products or add-on products installed. For example, IBM Data Server Runtime Client.

Do not use this type of fix pack if the installed DB2 database products are only DB2 database server products or a Data Server Client. Instead, use the single server image fix pack.

For Windows platforms, if you have more than one DB2 database product (which includes at least one product that is not a Data Server Client or a DB2 database server) installed in a single DB2 copy, you must download and uncompress all of the corresponding product-specific fix packs before starting the fix pack installation process.

- A universal fix pack.

The universal fix pack services installations where more than one DB2 database product has been installed.

The universal fix pack is not needed if the installed DB2 database products are only DB2 database server products or a Data Server Client. In this case, the single server image fix pack should be used.

On Linux or UNIX operating systems, if national languages have been installed, you also require a separate national language fix pack. The national language fix pack can not be installed alone. A universal or product-specific fix pack must be applied at the same time and they must both be at the same fix pack level. For example, if you are applying a universal fix pack to non-English DB2 database products on Linux or UNIX, you must apply both the universal fix pack and the national language fix pack to update the DB2 database products.

Restrictions

- A DB2 Version 10.1 fix pack can only be applied to DB2 Version 10.1 general availability (GA) or DB2 Version 10.1 fix pack copies.
- All DB2 instances, DAS, and applications related to the DB2 copy being updated must be stopped before installing a fix pack.
- In a partitioned database environment, before installing the fix pack, you must stop the database manager on all database partition servers. You must install the fix pack on the instance-owning database partition server and all other database partition servers. All computers participating in the instance must be updated to the same fix pack level.
- On Linux or UNIX operating systems:
 - If you have DB2 database products on a Network File System (NFS), you must ensure the following applications are stopped completely before installing the fix pack: all instances, the DB2 administration server (DAS), interprocess communications (IPC), and applications on other machines using the same NFS mounted installation.
 - If the system commands **fuser** or **lsof** are not available, the **installFixPack** command cannot detect loaded DB2 database files. You must ensure no DB2 files are loaded and provide an override option to install the fix pack. On UNIX, the **fuser** command is required to check for loaded files. On Linux, either the **fuser** command or **lsof** command is required.

For details on the override option, see the **installFixPack** command.

- On client applications, after a fix pack has been applied, to perform autobind of applications, the user must have bind authority.
- Installation of a DB2 fix pack will not service IBM Data Studio.

Procedure

To install a fix pack:

1. Check fix pack prerequisites.
2. Perform the necessary tasks before installing a fix pack.
3. Choose a fix pack installation method and install the fix pack.

4. Perform the necessary tasks after installing the fix pack.
5. Apply the appropriate DB2 database product license.

If a previously licensed copy of a DB2 database server product does not already exist on the machine, a single server fix pack image can be used to install any of the DB2 database server products. In this case, the DB2 database product installed is treated as a try and buy license, and will stop working after a 90 day trial period unless you upgrade the try and buy license.

What to do next

Check the log file for any post-installation steps, or error messages and recommended actions.

For non-root installations on Linux or UNIX, root-based features (such as High Availability and operating system-based authentication) can be enabled using the **db2rfe** command. If root-based features were enabled after installing your DB2 database product, you must rerun the **db2rfe** command each time a fix pack is applied in order to re-enable those features.

If you have multiple DB2 copies on the same system, those copies can be at different version and fix pack levels. If you want to apply a fix pack to one or more DB2 copies, you must install the fix pack on those DB2 copies one by one.

Chapter 52. Preparing to install a fix pack

In order to install a fix pack, you must first download and uncompress the fix pack. If you already installed DB2 for Linux, UNIX, and Windows in the selected path, you must also stop various DB2 processes.

Before you begin

If a IBM PowerHA SystemMirror for AIX cluster is running, you cannot perform a IBM Tivoli System Automation for Multiplatforms (SA MP) installation, upgrade, or update because SA MP bundles Reliable Scalable Cluster Technology (RSCT) file sets that are dependent on PowerHA SystemMirror. To skip the SA MP installation use the **db2setup** command or the **installFixPack** command. For information about installing or upgrading SA MP using a PowerHA SystemMirror cluster, see the white paper entitled “Upgrade guide for DB2 Servers in HACMP™ Environments”, which is available from the IBM Support and downloads website (<http://www.ibm.com/support/docview.wss?uid=swg21045033>).

Procedure

To prepare to install a fix pack:

1. Check the fix pack prerequisites. See “Checking fix pack prerequisites” on page 460.
2. Get the fix pack. See “Getting fix packs” on page 461.
3. Uncompress the fix pack.
 - For Linux and UNIX operating systems, see “Uncompressing fix packs (Linux and UNIX)” on page 462.
 - For Windows operating systems, see “Uncompressing fix packs (Windows)” on page 462.
4. Optional: Reduce the size of the fix pack.

You can use the **db2iprune** command to reduce the size of your DB2 fix pack installation image. Before installing a pruned fix pack, you must ensure that the pruned fix pack image contains at least the same components that are in the DB2 copy. If you prune too much from the fix pack image, the fix pack installation fails.
5. If you already have DB2 database products installed in the selected installation path:
 - a. Optional: Back up your current configuration and diagnostic information.

Gathering diagnostic information might be helpful when you are about to install a fix pack. This information will help to diagnose any problems that might arise after the installation. See “Backing up DB2 server configuration and diagnostic information” in *Upgrading to DB2 Version 10.1*.
 - b. Perform one of the following actions:
 - For Linux and UNIX operating systems, stop all DB2 processes. See “Stopping all DB2 processes (Linux and UNIX)” on page 463.
 - For Windows operating systems, stop all DB2 instances, services, and applications. See “Stopping all DB2 instances, services and applications (Windows)” on page 464.

What to do next

Choose a fix pack installation method. See Chapter 53, “Fix pack installation methods,” on page 467.

Checking fix pack prerequisites

Ensure that you meet all of the software, hardware and operating system prerequisites before you download a fix pack.

Procedure

Before downloading a fix pack, perform the following steps:

1. Ensure that your system meets all of the installation requirements.
2. Ensure that your system meets all of the installation requirements. Refer to Part 1, “Installation requirements for DB2 database products,” on page 1.
This action prevents technical problems that might occur after the installation and configuration of the fix pack.
3. For 64-bit DB2 for Linux on System i and pSeries®, ensure that the IBM XL C/C++ Compiler run time is installed.
Download and install the IBM XL C/C++ Run-Time Environment Component. Refer to: <http://www.ibm.com/support/docview.wss?uid=swg24007906>
4. For AIX, if DB2 database products already exist on the system and the Trusted Computing Base (TCB) is enabled, ensure it is in a consistent state.
Ensure that the TCB is in a consistent state for any file or directory related to a DB2 instance, the DB2 Administration Server, and installed DB2 files. The TCB can be in an inconsistent state if, for example, DB2 instances were removed without using the **db2idrop** command.
To check the TCB state, enter the following command as root:

```
tcbck -n ALL
```


Refer to AIX documentation for details regarding the Trusted Computing Base.
5. Review the Flashes and open Authorized Problem Analysis Reports (APARs) on the DB2 for Linux, UNIX and Windows product support website:
http://www.ibm.com/software/data/db2/support/db2_9/.
Open APARs describe defects that are not yet addressed in a fix pack, and they might contain workarounds that will impact your use of the database system. For a list of open DB2 for Linux, UNIX and Windows APARs, refer to:
<http://www.ibm.com/support>
6. Ensure that you have sufficient file system free space to download and extract the fix pack.
On Linux and UNIX, you need to have a file system with one gigabyte of free space to contain the .tar.gz file and the uncompressed installation image. If you also plan to install a national language fix pack, you need up to two gigabytes of free space.
7. Ensure that you have the free space required to install the fix pack.
The free space must be available in the location of the DB2 copy that you are updating or in the location where you plan to create a new DB2 installation.
 - On Windows, the space required to install the fix pack is equal to the space required for the DB2 Version 10.1 for Linux, UNIX and Windows general availability (GA) installation. Ensure the space required is available in the

location of the DB2 copy that you are updating or in the location where you plan to create a new installation of DB2. Refer to “Disk and memory requirements” on page 31.

- On Linux and UNIX:
 - If you do not already have a DB2 database product installed and you are using the fix pack to create a new installation, refer to “Disk and memory requirements” on page 31.
 - If you already have DB2 database product installed, the space required to install the fix pack is equal to the space consumed by the existing DB2 database products. This space is only required temporarily during the fix pack installation process.

To determine the space used by the existing DB2 database products, perform the command:

```
du -k -s DB2DIR
```

where *DB2DIR* represents the location where the DB2 database product is installed.

8. If you already have a DB2 database product installed and have obtained special fixes, contact IBM support to determine whether you need an updated version of the fixes before you install a fix pack.

This helps to ensure that your system is in a consistent state and that no special fixes are lost.

9. Optional: Review the “Product overviews” in the *DB2 Information Center*.

Results

Once you have ensured that you meet all of these prerequisites, complete the remainder of the necessary tasks before installing a fix pack.

Getting fix packs

After checking the fix pack prerequisites, download the fix pack.

Before you begin

Check the fix pack prerequisites. See “Checking fix pack prerequisites” on page 460.

Procedure

To get a fix pack:

1. Determine which fix pack you need.

In general, choose the most recent fix pack to avoid encountering problems caused by software defects that are already known and corrected.
2. Locate the fix pack on the DB2 for Linux, UNIX, and Windows product support website: www.ibm.com/support/docview.wss?rs=71&uid=swg27007053.

Ensure that you choose the appropriate fix pack for your operating system. Choose between DB2 database product-specific fix packs and universal fix packs.
3. Download the fix pack.

In most cases, you can choose to access the FTP folder directly or use a Java applet called Download Director to download the files.

On Linux and UNIX operating systems, there must not be any spaces in the directory path where you plan to download and uncompress the fix pack. If there are spaces in the directory path, the installation fails. For example, make sure that your directory path resembles the following example:
/home/DB2FixPack/FP1/. It must not resemble the following: /home/DB2
FixPack/FP1/.

What to do next

After you successfully download the fix pack, perform the remainder of the preparatory steps before installing a fix pack. See Chapter 52, “Preparing to install a fix pack,” on page 459.

Uncompressing fix packs (Linux and UNIX)

All fix pack installation images on the FTP site are compressed using **gzip**. Before you can install a fix pack, you must copy the image to a temporary directory and use **gunzip** and **tar** to extract the fix pack installation image.

Procedure

To uncompress a fix pack installation image, perform the following steps:

1. Copy the gzipped image to a temporary location.
2. Change to the directory where you copied the image.
3. Enter the following command to uncompress the file:

```
gunzip -c filename.tar.gz | tar -xvf -
```

where *filename* is the fix pack you are installing.

Note: **gunzip** is part of the AIX 5L™ default installation setup. If you do not have **gunzip**, install the rpm.rte fileset from the AIX 5L installation media. The rpm.rte fileset contains **gunzip**. You can also download **gzip** for AIX 5L from the website: www.ibm.com/servers/aix/products/aixos/linux/rpmgroups.html

Uncompressing fix packs (Windows)

Fix pack installation images on the FTP site are in compressed format. Before you can install a fix pack, you must uncompress the fix pack installation image in a temporary directory.

Before you begin

- If you have more than one DB2 database product installed in a single DB2 copy, you must download and uncompress all of the corresponding product-specific fix packs.
- If you plan to install the fix pack in silent mode, all of the fix pack images must be uncompressed into subdirectories under the same parent directory.
- If you plan to install the fix pack using the DB2 Setup wizard, the **setup** command detects the other installed DB2 database products. If all of the fix pack images for your DB2 database products are uncompressed into subdirectories under the same parent directory, the **setup** command automatically starts installing the fix pack for all DB2 database products without prompting you. If the fix pack product images are uncompressed into different directories, the **setup** command detects the other installed DB2 database products and prompts you for each directory path.

Procedure

To uncompress a fix pack installation image, perform the following steps:

1. Change to the directory where you have downloaded the fix pack installation image.
2. Double click the self-extracting .exe file. The WinZip Self Extractor window opens.
3. Select a folder to contain the unzipped files.
4. Click Unzip. All files are unzipped to the specified folder.

What to do next

Repeat these steps to uncompress the self-extracting images for all of the DB2 database products you have installed. If you want to perform a silent installation using a response file, all of the images must be uncompressed to subdirectories under the same parent directory.

Stopping all DB2 processes (Linux and UNIX)

Before installing a fix pack, if there are DB2 database products installed in the selected installation path, you must stop all of the DB2 processes. If you have multiple DB2 copies, stop only the DB2 processes that are associated with the copy that you are updating.

Procedure

To stop all DB2 processes, perform the following steps:

1. Log on as root.
2. Determine which instances are associated with the DB2 copy. Issue the command:

```
DB2DIR/instance/db2ilist
```

where *DB2DIR* represents the location where the DB2 copy is installed.

3. Run the following commands for each instance in the DB2 copy:

```
su - iname
. $HOME/sqllib/db2profile
db2 force applications all
db2 terminate
db2stop
db2licd -end      # run at each physical partition
exit
```

where *iname* represents the instance owner name. If you are an PowerHA SystemMirror user, you must use the **ha_db2stop** command to stop DB2 instead of the **db2stop** command. If you use the **db2stop** command instead of the **ha_db2stop** command, you will trigger a failure event.

4. If the DB2 Administration Server (DAS) belongs to the DB2 copy that you are updating, stop the DAS:

```
su - aname
. $HOME/das/dasprofile
db2admin stop
exit
```

where *aname* represents the DAS owner name.

Note: Since there can only be one DAS on the system, this step affects all other DB2 copies on the system.

5. Optional: On AIX, run **slibclean** to unload unused shared libraries from memory before installation:

```
/usr/sbin/slibclean
```

6. Disable the fault monitor processes. To stop the Fault Monitor Daemon, issue the command:

```
DB2DIR/bin/db2fm -i iname -D
```

where *DB2DIR* is the location where the DB2 copy is installed and *iname* represents the instance owner name. The command must be performed once for each instance in the DB2 copy.

7. If the Fault Monitor Coordinator (FMC) is started, prevent your instances from auto-starting:

- a. To determine whether the FMC is started, issue the command:

```
DB2DIR/bin/db2fmcu
```

where *DB2DIR* is the location where the DB2 copy is installed. If the FMC is started, you will see output similar to the following: FMC: up: PID = 3415 . If the FMC is disabled, the output from the **db2fmcu** command will be: FMC: down.

- b. If the FMC is started, determine whether any instances are configured to auto-start after each system restart. Issue the command:

```
DB2DIR/instance/db2iset -i iname -all
```

where *DB2DIR* is the location where the DB2 copy is installed and *iname* represents the instance owner name. The command must be performed once for each instance in the DB2 copy. If the output from the **db2iset** command includes the following, it means that the instance is configured to auto-start: DB2AUTOSTART=YES

- c. Prevent the instances from auto-starting. Issue the command:

```
DB2DIR/instance/db2iauto -off iname
```

where *DB2DIR* is the location where the DB2 copy is installed and *iname* represents the instance owner name. After you have completed the fix pack installation, you can re-start instance auto-start:

```
DB2DIR/instance/db2iauto -on iname
```

8. Ensure all DB2 interprocess communications are cleaned for the instance to be updated. As the instance owner, run the following command at each physical partition:

```
$HOME/sql1lib/bin/ipclean
```

Stopping all DB2 instances, services and applications (Windows)

Before installing a fix pack, if there are DB2 database products installed in the selected installation path, you must stop all DB2 instances, all DB2 services, and all applications. If you have multiple DB2 copies, stop only the DB2 processes that are associated with the copy that you are updating.

Procedure

To stop the appropriate instances, services and applications:

1. Determine which instances are associated with the DB2 copy by issuing the command:

```
DB2PATH\bin\db2ilist
```

where *DB2PATH* represents the location where the DB2 copy is installed.

2. Stop all instances and DB2 services, using the services control panel applet: **Control Panel > Administrative Tools > Services**. If you have active database clients, force these clients off when stopping the instance. To force clients, issue the following command:

```
db2stop force
```

Alternatively, use the task assistant for stopping instances and forcing applications or client off the DB2 server in IBM Data Studio.

Chapter 53. Fix pack installation methods

There are several different methods for installing fix packs, depending on the operating system and whether or not there are existing database products.

On Linux and UNIX operating systems, there are different methods that you can use the fix pack and the installation methods match these purposes:

- To update existing DB2 database products, with an ese or dsf instance type, refer to “Installing a fix pack to update a Version 10.1 DB2 pureScale instance” on page 469.

Follow these instructions if a DB2 database product is already installed and you want to apply a new fix pack level. The **installFixPack** command is used to install the fix pack.

- To install new DB2 database products in a selected location, refer to “Installing a fix pack to install new DB2 database products (Linux and UNIX)” on page 473.

Follow these instructions to install new DB2 database products at a particular fix pack level. The **db2setup** command is used to perform the installation.

- To update existing DB2 database products in a selected location, refer to “Installing a fix pack to update existing DB2 database products (Linux and UNIX)” on page 473.

Follow these instructions if a DB2 database product is already installed and you want to apply a new fix pack level. The **installFixPack** command is used to install the fix pack.

On Windows operating systems, the same commands are used whether you use the fix pack to update existing DB2 database products or add new DB2 database products. The installation methods differ according to the number of DB2 database products involved and whether DB2 is configured to use Microsoft Cluster Server (MSCS), as follows:

- “Installing a fix pack for a single database product (Windows)” on page 475.

Follow these instructions if a single DB2 database product is already installed and you want to apply a new fix pack level. The **setup** command is used to install the fix pack.

- “Installing a fix pack for multiple database products (Windows)” on page 475.

Follow these instructions to install a fix pack on a system with multiple DB2 database products installed. The **setup** command is used to install the fix pack.

- “Installing a fix pack using a response file (Windows)” on page 476 (single product or multiple products).

Follow these instructions to perform a response file installation of a fix pack. A response file installation might also be referred to as a silent installation or an unattended installation. The **setup** command is used to install the fix pack.

- “Installing a fix pack in a Microsoft Cluster Server environment (Windows)” on page 477.

Follow these instructions to install a fix pack on top of an existing DB2 installation in a MSCS environment.

Universal versus product-specific fix packs

You have two choices when it comes to fix pack installation; using a universal fix pack (all products) or a product-specific fix pack.

Universal [all products] fix pack

Similar to previous versions of DB2 database products, you can use the two-step process of installing a GA level product then applying the desired fix pack level. For Windows operating systems, you can use the universal image to install DB2 to a new location. This installation method applies to platforms where you have one product installed in your install path. For example, you might have an existing DB2 database product installation, and now want to upgrade in place to the Fix Pack 1 level using steps such as:

1. Download desired universal fix pack (if you have only one product installed, you can also use the product-specific fix pack, see following section).
2. Stop all instances associated with this copy.
3. Stop the DAS if it is associated with this copy.
4. Apply the fix pack.
5. Restart the instances (**db2start**), including the DAS (**db2admin start**) if it is associated with the current copy.

Product-specific fix pack

Use the product-specific fix pack to upgrade a single product, or to install to a new path. If you choose to install the product-specific fix pack, you can download the products at the desired level and then install the products in a single step. For example, if you want to install Enterprise Server Edition (ESE) at the Fix Pack 1 level, perform the following steps:

1. Download the product at the desired level.
2. Install the product to a new path.
3. Test the new installation.
4. Drop the test instances.
5. Run **db2iupdt** from the new install copy to move your instances over to the new fix pack level.

This approach allows you to test the new version on your production system without impacting your production databases. You now have a 15-30 second window of downtime (the time it takes to issue the **db2stop**, **db2 bind**, and **db2start** commands).

Product-specific fix packs such as Spatial Extender, and InfoSphere Federation Server are available for cases where you are adding this functionality to an already upgraded DB2 copy. If you have ESE at Fix Pack 1, you need Spatial Extender at Fix Pack 1 to install over top of the existing ESE product. The Spatial Extender level on your GA DVD is no longer sufficient. In this case, you must download the Spatial Extender Fix Pack 1 image, add the license file, run the installer, and, when the installer is finished, you will have a working Spatial Extender installation.

Upgrading multiple products installed in the same install path

In cases where you have more than one product installed to one install path, for example you have ESE and QP installed, you can either use one ESE image and one QP image to install to a new location, or you use the universal fix pack to upgrade in place. You need to know which approach you want to use before downloading. If you only have one product installed, for example ESE, then the ESE product-specific fix pack or the universal fix pack will both update in place.

Why would I use the universal fix pack?

- If I want to upgrade multiple products in place.

- If I do not remember what I have installed to be able to install the fix pack in a new location without losing my current set of features or products.

Why would I use the product-specific fix pack?

- If I have a single product and want to save the time to download the fix pack.
- If system downtime is costly.
- If I need to test the new fix pack without impacting my production databases.
- If I need the ability to go back to the previous version quickly (system downtime in case of failure).
- If I already have a fix pack installed and need to add new features or products to the fix pack.

Note: Many of these advantages can be achieved by installing from your original media to a new location and using the universal fix pack to upgrade features or products, however, there is a cost of extra installation time. If you have already downloaded the universal fix pack, this might take less time than downloading all the pieces you need to replicate your existing copy at the desired fix pack level.

Installing a fix pack (Linux and UNIX)

You can update an existing DB2 database product to a new fix pack level or install a new DB2 database product at a particular fix pack level.

Installing a fix pack to update a Version 10.1 DB2 pureScale instance

If a DB2 pureScale instance is already installed, follow these instructions to apply a new fix pack level.

Before you begin

- Ensure that you meet all of the necessary tasks before installing a fix pack.
- Ensure that you are logged in as root.
- The entire cluster must be offline.

Procedure

To update an existing DB2 pureScale instance with a fix pack:

1. Confirm the level of IBM Tivoli System Automation for Multiplatforms (SA MP) Base Component installed:
 - a. Query the SA MP Base Component level already installed on your system:
`<image_path>/db2/<platform>/tsamp/db2cktsa -v install`
 - b. Query the level of SA MP Base Component on the installation media:
`<image_path>/db2/<platform>/tsamp/db2cktsa -v media`
2. Confirm the level of IBM General Parallel File System (GPFS) installed:
 - a. Query the level if GPFS already installed on your system:
`<image_path>/db2/<platform>/gpfs/db2ckgpfs -v install`
 - b. Query the level of GPFS on the installation media:
`<image_path>/db2/<platform>/gpfs/db2ckgpfs -v media`

3. If a member or CF is running on the host, the member must be stopped:
`db2stop member <member-id>`
`db2stop cf <cf-id>`

4. Stop the DB2 instance on all hosts by issuing the following command on a single host:

```
su - InstName
db2stop
exit
```

where *InstName* represents the instance owner name.

5. If either the level of SA MP Base Component or the level of GPFS software on the installation media is higher than the level already installed on your system, the DB2 cluster services administrator must place all hosts in the cluster into maintenance mode.

- a. Stop the instance on each host:

```
db2stop instance on hostname
```

where *hostname* represents the host name for a given member or CF.

- b. To place the host into maintenance mode, enter the **db2cluster** command on one host in the DB2 pureScale environment:

```
DB2DIR/bin/db2cluster -cm -enter -maintenance -all
```

where *DB2DIR* represents the installation location of your DB2 copy.

6. If the level of GPFS on the installation media is higher than the level already installed on your system, the DB2 cluster services administrator must place the host into maintenance mode. To place the host into maintenance mode, enter the **db2cluster** command on one host in the DB2 pureScale environment:

```
DB2DIR/bin/db2cluster -cfs -enter -maintenance -all
```

where *DB2DIR* represents the installation location of your DB2 copy.

7. Update each host in the DB2 pureScale instance to Version 10.1 Fix Pack 2 (or higher) on each host. For example, to update to Fix Pack 2, run the **installFixPack** command on each host:

```
installFixPack -b /opt/fp1level/path -f 1e1e1 -p /opt/fp2level/path -t /tmp/trace-file-path -L -l /tmp/install-log-path
```

The path specified in the **-p** parameter must be different than the path specified in the **-b** parameter.

8. If the DB2 cluster services administrator placed the SA MP Base Component host into maintenance mode, the DB2 cluster services administrator must remove the host from maintenance mode. To remove the SA MP Base Component host from maintenance mode, enter the **db2cluster** command on one host in the DB2 pureScale environment:

```
DB2DIR/bin/db2cluster -cm -exit -maintenance -all
```

where *DB2DIR* represents the installation location of your DB2 copy.

9. If the DB2 cluster services administrator placed the GPFS host into maintenance mode, the DB2 cluster services administrator must remove the host from maintenance mode. To remove the GPFS host from maintenance mode, enter the **db2cluster** command on one host in the DB2 pureScale environment:

```
DB2DIR/bin/db2cluster -cfs -exit -maintenance -all
```

where *DB2DIR* represents the installation location of your DB2 copy.

10. On each host, update the instance to the new installation path:

```
DB2DIR/instance/db2iupdt <InstName>
```

where *DB2DIR* represents the new fix pack directory specified in a previous step, and *InstName* specifies the name of the instance.

In a geographically dispersed DB2 pureScale cluster (GDPC) environment, the **db2iupdt** command issues a warning indicating that the **db2cluster** command failed (**db2cluster -cfs -verify -configurations**). To avoid this warning, before running the **db2iupdt** command, enter:

```
root@hostA1>export PASS_USER_MANAGED_GPFS_VALIDATION=YES
```

11. The DB2 cluster services administrator must commit the updates made to DB2 cluster services and make them available to the DB2 database system. On one host, run the **db2cluster** commands:

```
DB2DIR/bin/db2cluster -cfs -commit
DB2DIR/bin/db2cluster -cm -commit
```

where *DB2DIR* represents the new fix pack level installation location of your DB2 copy.

12. Determine the host failure detection time for the current RSCT communication group. To retrieve the detection time:

```
db2cluster -cm -list -HostFailureDetectionTime
```

13. If you are planning to use multiple communication adapter ports, install and configure additional adapters, switches and cables.

14. Configure RSCT communication group host failure detection time for the new cluster interconnect netname. As root, change the host failure detection time:

```
db2cluster -cm -set -option HostFailureDetectionTime -value <value>
```

where *value* is the detection time retrieved previously.

15. If you are planning to use multiple communication adapter ports, update the member or CF servers to use multiple netnames:

```
db2iupdt -update -cf <CFHostName> -cfnet CFnetname1,CFnetname2 <InstName>
db2iupdt -update -m <MemberHostName> -mnet MemberNetname1,MemberNetname2 <InstName>
```

16. As the instance owner, refresh the resource model:

```
db2cluster -cm -repair -resources
```

17. If the hosts were placed into maintenance mode, the DB2 cluster services administrator must start the DB2 instance on all hosts. To start the DB2 instance run the **db2start** command on each host in the DB2 pureScale environment

```
db2start instance on <hostname>
```

where *<hostname>* represents the host name for a given member or CF, and the **db2start** instance on *hostname* command is run for each host in the cluster.

18. To start the database manager, run the **db2start** command:

```
su - <iname>
db2start
exit
```

where *<iname>* represents the instance owner name.

19. To uninstall the old DB2 copy, run the **db2_deinstall** command:

```
DB2DIR/install/db2_deinstall -a
```

where *DB2DIR* is the installation path of the old DB2 copy.

Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP) Base Component in a DB2 pureScale environment

You can upgrade IBM Tivoli System Automation for Multiplatforms (SA MP) using either the DB2 fix pack installer or the **installSAM** install script that is included in the IBM DB2 pureScale Feature for Enterprise Server Edition install media.

Before you begin

Whether you are using the DB2 fix pack installer (**installFixPack** command), **installSAM**, or **uninstallSAM**, you must meet the following basic prerequisites for installing, upgrading, or uninstalling SA MP:

- To install and use SA MP, your system configuration and intended use of SA MP must meet the terms of the license that comes with the SA MP that is integrated with the DB2 pureScale Feature
- To install or upgrade SA MP your system architecture must be supported by the SA MP that is integrated with the DB2 pureScale Feature
- You must have root authority to install SA MP.
- the peer domain must be offline or in maintenance mode on the current node

To upgrade SA MP using **installFixPack**, perform the following steps before installation:

1. Go to the Passport Advantage website and obtain a permanent SA MP license file (**sam32.lic** or **sam31.lic**) from your entitled DB2 Version 9.8 activation CD.
2. Copy the permanent license file into the *Fixpack-path/db2/platform/tsamp/license* directory where *Fixpack-path* represents the path where the fix pack image is located and *platform* represents the operating system used.
3. Remove the **sam31tb.lic** or the **sam32tb.lic** file from the fix pack image. The installation will fail if you do not delete these additional license files.
4. Continue with the upgrade process.

Procedure

To upgrade SA MP in a DB2 pureScale environment, first update the TSA binaries, then upgrade the peer domain:

1. Update the TSA binaries by running either the **installFixPack**, or running the **installSAM** command:
 - a. Run the **installFixPack** command
 - b. To run the **installSAM** install script: The **installSAM** install script is located on the DB2 pureScale Feature for Enterprise Server Edition install media at the following location:
`db2/platform/tsamp`

where *platform* refers to the appropriate hardware platform. For information about using **installSAM** see: IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide Version 3.2 or later

2. To upgrade the Reliable Scalable Cluster Technology (RSCT) peer domain:
You can perform either a node-by-node migration or migrate the entire domain in 1 step. For more information, see the IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide Version 3.2 or later (SC34-2584-01) and go to Chapter 1 “Installing on UNIX and Linux”, section

“Installing and upgrading IBM Tivoli System Automation”. Complete the steps in the “Migrating an entire domain” section.

Installing a fix pack to update existing DB2 database products (Linux and UNIX)

If a DB2 database product is already installed, follow these instructions to apply a new fix pack level.

Before you begin

- Ensure that you meet all of the necessary tasks before installing a fix pack. Refer to Chapter 52, “Preparing to install a fix pack,” on page 459.
- If there is more than one DB2 database product installed in the selected path, you must use a universal fix pack image to install the fix pack.
- If you want to update an existing DB2 database product that has national languages installed, you must obtain the national language fix pack in addition to either an individual fix pack or a universal fix pack. National language fix packs can not be used alone.

For example, to install a fix pack on a DB2 database product with non-English support installed, download the DB2 database product-specific fix pack image (or the universal fix pack image) and the national language fix pack. Then run **installFixPack** from the DB2 database product-specific fix pack image (or the universal fix pack image).

- To install a fix pack to a DB2 pureScale instance, follow the instructions in “Installing a fix pack to update a Version 10.1 DB2 pureScale instance” on page 469.

Procedure

To install a fix pack:

1. For root installations, log on as root. For non-root installations, log on with the user ID that owns the non-root installation.
2. Change to the directory that contains the fix pack image.
3. Launch the installation by issuing the **installFixPack** command. For example,

```
./installFixPack -b DB2DIR
```

where *DB2DIR* is the location of the DB2 database products that you want to update.

In clustered environments where some instances are not mounted, add the **-f ha_standby_ignore** option. For example,

```
./installFixPack -b DB2DIR -f ha_standby_ignore
```

What to do next

To complete the installation, perform the necessary post-installation tasks for fix packs. Refer to “Post-installation tasks for fix packs (Linux and UNIX)” on page 482.

Installing a fix pack to install new DB2 database products (Linux and UNIX)

Follow these instructions to install new DB2 database products at a particular fix pack level. The **db2setup** command is used to perform the installation.

Before you begin

- Perform all of the necessary tasks before installing a fix pack. Refer to Chapter 52, “Preparing to install a fix pack,” on page 459.
- If you want to install a new DB2 database product with national languages enabled, you must obtain the national language fix pack in addition to either an individual fix pack or a universal fix pack. National language fix packs can not be used alone.

For example, to install a DB2 database product with non-English support, download the DB2 database product-specific fix pack image and the national language fix pack. Then run **db2setup** from the DB2 database product-specific fix pack image.

- The **db2setup** command is located in the single server fix pack image, not in the universal fix pack image.

About this task

Use this method if:

- There are no DB2 database products in the selected installation path, or
- DB2 database products exist in the selected installation path and you want to add more products (at the same fix pack level as the existing products) in the same path.

Procedure

To install the DB2 database products:

1. Log on as root.
2. Change to the directory that contains the fix pack image.
3. Launch the installation by issuing the command:

```
./db2setup
```

Note:

- If you choose to install InfoSphere Federation Server products, you will receive a warning message if you issue **db2setup** from a fix pack image that is at a lower level than the existing DB2 database products. You must apply the appropriate fix pack after the product installation is complete.
- For all other DB2 database products, if the fix pack image that you are using to add new products is not at the same level as the installed DB2 database products in a particular DB2 copy, the installation attempt will not succeed. For example:
 - If the fix pack image is at a lower level than the installed products, the error message returned by **db2setup** will advise you to obtain the right fix pack image.
 - If the fix pack image is at a higher level than the installed DB2 database products, the error message returned by **db2setup** will advise you to upgrade the existing products using **installFixPack** first.

What to do next

To complete the installation, perform the necessary post-installation tasks for fix packs. Refer to “Post-installation tasks for fix packs (Linux and UNIX)” on page 482.

Installing a fix pack (Windows)

You can install a fix pack for a single database product or for multiple database products.

You can also install a fix pack using a response file or in a Microsoft Cluster Server (MSCS) environment.

Installing a fix pack for a single database product (Windows)

Follow these instructions if a single DB2 database product is already installed and you want to apply a new fix pack level. The **setup** command is used to install the fix pack.

Before you begin

- Ensure that you meet all of the fix pack prerequisites.
- Ensure that you have the appropriate installation user accounts. In general, the user account must belong to the Administrators group on the machine where you will perform the installation.

Procedure

To install a fix pack:

1. Change to the folder where the unzipped files are located. The **setup** command is located under the folder labelled with the abbreviated product name. For example, DB2 Enterprise Server Edition is under ESE.
2. Double click the setup.exe file to start the DB2 Setup wizard. The DB2 Setup wizard launchpad opens. Online help is available to guide you through the wizard. To invoke online help, click **Help** or press F1.

What to do next

To complete the installation, perform the necessary post-installation tasks for fix packs.

Installing a fix pack for multiple database products (Windows)

Follow these instructions to install a fix pack on a system with multiple DB2 database products installed. The **setup** command is used to install the fix pack.

Before you begin

- Ensure that you meet all of the fix pack prerequisites. Refer to “Checking fix pack prerequisites” on page 460.
- Ensure that you have the appropriate installation user accounts. In general, the user account must belong to the Administrators group on the machine where you will perform the installation.
- If there is more than one DB2 database product installed in the selected path, you can use a universal fix pack image to install the fix pack.

Procedure

To install a fix pack:

1. Change to the folder where the unzipped files are located. The **setup** command is located under the folder labelled with the abbreviated product name. For example, DB2 Enterprise Server Edition is under ESE.

2. Double click the `setup.exe` file to start the DB2 Setup wizard. The DB2 Setup wizard launchpad opens.

The DB2 Setup wizard detects the installed DB2 database products.

- If all the product images are uncompressed into subdirectories under the same parent directory, the DB2 Setup wizard automatically starts the installation of all DB2 database products without prompting you.
- If the product images are uncompressed into different directories, the DB2 Setup wizard detects the installed DB2 database products and prompts you for the directory paths.

Online help is available to guide you through the wizard. To start online help, click **Help** or press F1.

What to do next

Perform the necessary post-installation tasks for fix packs. Refer to “Post-installation tasks for fix packs (Windows)” on page 481.

Installing a fix pack using a response file (Windows)

Follow these instructions to perform a response file installation of a fix pack. A response file installation might also be referred to as a silent installation or an unattended installation. The **setup** command is used to install the fix pack.

Before you begin

- Ensure that you meet all of the necessary tasks before installing a fix pack. Refer to Chapter 52, “Preparing to install a fix pack,” on page 459.
- Ensure that you have the appropriate installation user accounts. In general, the user account must belong to the Administrators group on the machine where you will perform the installation.

Procedure

To install a fix pack using a response file:

1. Change to the folder where the unzipped files are located.

The **setup** command is located under the folder labelled with the abbreviated product name. For example, DB2 Enterprise Server Edition is under ESE.

2. Ensure that all of the fix pack installation images were uncompressed to subdirectories under the same parent directory.

If you have more than one DB2 database product installed, the DB2 Setup wizard detects the other installed products. The installation fails if all of the appropriate DB2 database products' uncompressed fix pack images are not in subdirectories under the same parent directory.

3. Install the fix pack using the **setup** command with the response file (**-u**) option. For example, enter:

```
setup -u c:\db2fixpk.rsp -t c:\db2fixpk.trc -l c:\db2fixpk.log
```

where `db2fixpk.rsp` is the response file name and **-t** and **-l** are optional parameters that specify a trace file and a log file, respectively. Sample response files can be found in the *product-abbreviation*\db2\Windows\samples directory in the fix pack installation image. For example, ESE\db2\windows\samples.

What to do next

To complete the installation, perform the necessary post-installation tasks for fix packs. Refer to “Post-installation tasks for fix packs (Windows)” on page 481.

Installing a fix pack in a Microsoft Cluster Server environment (Windows)

Follow these instructions to install a fix pack on top of an existing DB2 installation in a Microsoft Cluster Server (MSCS) environment.

Before you begin

- Ensure that you meet all of the necessary tasks before installing a fix pack. Refer to Chapter 52, “Preparing to install a fix pack,” on page 459.
- Ensure that you have the appropriate installation user accounts. In general, the user account must belong to the Administrators group on the machine where you will perform the installation.

About this task

To illustrate how to install a fix pack in an MSCS environment, a simple example configuration must be understood. In this example, the initial configuration is a DB2 instance that is composed of two database partitions. In addition, there are two Cluster Groups:

- DB2 Group 0: contains Partition 0 active on Machine A. The DB2 Administration Server (DAS) is also located here.
- DB2 Group 1: contains Partition 1 active on Machine B.

This is the initial configuration.

Procedure

To install a fix pack on top of an existing DB2 installation in an MSCS environment:

1. Set automatic failback to off.

During the installation process you might have to restart your machine. When this happens the cluster service restarts automatically. Therefore the automatic failback must be set off so none of the groups automatically failback.

For example, to disable automatic failback in DB2 Group 0:

- a. From the Cluster Administrator window, right-click **DB2 Group 0**.
- b. Select **Properties**. The DB2 Group 0 Properties window opens.
- c. On the **Failback** tab, select the **Prevent failback** radio button.
- d. Click **OK**.

Repeat these steps to disable automatic failback in DB2 Group 1.

2. Decide which machine you will install the fix pack on first.

In this example, Machine B will be upgraded first.

3. Move the DB2 Group off Machine B.

For example, to move DB2 Group 1 from Machine B to Machine A:

- a. From the Cluster Administrator window, click **DB2 Group 1**.
- b. Right-click, select **Move Group**. The Owner column changes to Machine A.

4. Stop the cluster server on Machine B.

For example:

- a. From the Component Services window, right-click **Cluster Service**.
 - b. Click **Stop**.
5. If there are multiple DB2 database products installed at the same location (in the same DB2 copy) on Machine B, install the fix pack for multiple database products. Otherwise, install the fix pack for a single database product. Refer to “Installing a fix pack for multiple database products (Windows)” on page 475 or “Installing a fix pack for a single database product (Windows)” on page 475, as appropriate.

Note:

- DB2 is still running and available on Machine A.
 - As part of the installation process, you might have to restart your computer.
 - The silent mode installation process can optionally be used to install the fix pack using a response file. Refer to “Installing a fix pack using a response file (Windows)” on page 476.
 - If db2sysstray.exe attempts to access an offline instance, then you might receive error message SQL5005C near the end of installation. This does not indicate an installation failure.
6. Take the DB2 resources offline.
- In the example, Partition 0, Partition 1, and the DAS are on Machine A at this point. These must be taken offline one at a time. For example:
- a. From the Cluster Administrator window, in the left window panel, select **Groups**.
 - b. Select **DB2 Group 0**. The resources for the group display in the right window panel.
For DB2 Group 0, the DB2 resources include Partition 0 and the DAS.
 - c. In the right window panel, right-click one of the resources. Select **Take Offline**.
Repeat this step for each of the DB2 resources in DB2 Group 0.
 - d. Select **DB2 Group 1**. The resources for the group display in the right window panel.
For DB2 Group 1, the DB2 resources include Partition 1.
 - e. In the right window panel, right-click the resource (Partition 1). Select **Take Offline**.
7. (Optional) Restart the cluster service on Machine B.
- If the fix pack installation required that you restart your computer, and if the **Restart Type** for the cluster service is set to automatic, skip this step because the cluster service is already started.
- To start the cluster service:
- a. From the Component Services window, right-click **Cluster Service**.
 - b. Click **Start**.
8. Move the DB2 Groups to Machine B.
- For example, to move DB2 Group 0 and DB2 Group 1 to Machine B:
- a. From the Cluster Administrator window, click **DB2 Group 0**.
 - b. Right-click, select **Move Group**. The Owner column changes to Machine B.
 - c. Click on **DB2 Group 1**.
 - d. Right-click, select **Move Group**. The Owner column changes to Machine B.
9. Bring the DB2 resources online.

In the example, Partition 0, Partition 1, and the DAS are on Machine B at this point. These must be brought back online one at a time. For example:

- a. From the Cluster Administrator window, in the left window panel, select **Groups**.
 - b. Select **DB2 Group 0**. The resources for the group display in the right window panel.
For DB2 Group 0, the DB2 resources include Partition 0 and the DAS.
 - c. In the right window panel, right-click one of the resources. Select **Bring Online**.
Repeat this step for each of the DB2 resources in DB2 Group 0.
 - d. Select **DB2 Group 1**. The resources for the group display in the right window panel.
For DB2 Group 1, the DB2 resources include Partition 1.
 - e. In the right window panel, right-click the resource (Partition 1). Select **Bring Online**.
10. Stop the cluster service on Machine A.
For example:
- a. From the Component Services window, right-click **Cluster Service**.
 - b. Click **Stop**.
11. If there are multiple DB2 database products installed at the same location (in the same DB2 copy) on Machine A, install the fix pack for multiple database products. Otherwise, install the fix pack for a single database product. Refer to “Installing a fix pack for multiple database products (Windows)” on page 475 or “Installing a fix pack for a single database product (Windows)” on page 475, as appropriate.

Note:

- DB2 is still running and available on Machine B.
 - As part of the installation process, you might have to restart your computer.
 - The silent mode installation process can optionally be used to install the fix pack using a response file. Refer to “Installing a fix pack using a response file (Windows)” on page 476.
 - If db2sysstray.exe attempts to access an offline instance, then you might receive error message SQL5005C near the end of installation. This does not indicate an installation failure.
12. (Optional) Restart the cluster service on Machine A.
If the fix pack installation required that you restart your computer, and if the **Restart Type** for the cluster service is set to automatic, skip this step because the cluster service is already started.
To start the cluster service:
- a. From the Component Services window, right-click **Cluster Service**.
 - b. Click **Start**.
13. Move the DB2 Groups back to the appropriate machines.
For example, to move DB2 Group 0 back to Machine A:
- a. From the Cluster Administrator window, click **DB2 Group 0**.
 - b. Right-click, select **Move Group**. The Owners column changes to Machine A.

This brings the MSCS environment back to the initial configuration.

14. If automatic failback was turned off in Step 1, turn it back on.

For example, to start automatic failback in DB2 Group 0:

- a. From the Cluster Administrator window, right-click **DB2 Group 0**.
- b. Select **Properties**. The DB2 Group 0 Properties window opens.
- c. On the **Failback** tab, select the **Allow failback** radio button.
- d. Click **OK**.

Repeat these steps to start automatic failback in DB2 Group 1.

What to do next

To complete the installation, perform the necessary post-installation tasks for fix packs. Refer to “Post-installation tasks for fix packs (Windows)” on page 481.

Chapter 54. After installing a fix pack

There is a set of tasks that can be performed after you install a fix pack. The tasks differ for Windows and Linux or UNIX operating systems.

Post-installation tasks for fix packs (Windows)

As part of a fix pack installation, binding of the database utilities (**IMPORT**, **EXPORT**, **REORG**, the Command Line Processor) and the CLI bind files occurs automatically.

However, if an error occurs, you can manually bind the database utilities and the CLI bind files. For partitioned database environments on certain Windows operating systems, you must start the DB2 Remote Command Service security feature. Recompiling applications is an optional task.

Procedure

Perform the following actions:

1. For partitioned database environments on Windows 2000 or later, start the DB2 Remote Command Service security feature to protect your data and resources.

To be fully secure, start either the computer (if the service is running under the context of the LocalSystem account) or a user for delegation (if the service is being run under the logon context of a user).

To start the DB2 Remote Command Service security feature:

- a. Open the Active Directory Users and Computers window on the domain controller, click **Start** and select **Programs > Administrative tools > Active Directory Users and Computers**
- b. In the right window panel, right-click the computer or user to start, select **Properties**
- c. Click the **General** tab and select the **Trust computer for delegation** check box. For user setting, click the **Account** tab and select the **Account is trusted for delegation** check box in the **Account option** group. Ensure that the **Account is sensitive and cannot be delegated** box has not been checked.
- d. Click **OK** to start the computer or user for delegation.

Repeat these steps for each computer or user that needs to be started. You must restart your computer for the security change to take effect.

If you want to disable the DB2 Remote Command Service security feature, enter the following command:

```
db2set DB2RCMD_LEGACY_MODE=ON
```

2. Optional: Update the system catalog objects in your databases to support the fix pack.

This task is strongly recommended if you want to use capabilities specific to the fix pack. This task is not necessary if you installed the fix pack to create a new installation, since there are no existing databases. For each instance in the DB2 copy where you applied the fix pack, perform the following actions:

- a. Determine which instances are associated with the DB2 copy by issuing the command:

```
DB2DIR\bin\db2ilist
```

where *DB2DIR* represents the location where the DB2 copy is installed.

- b. Perform the following command once for each database in the instances:

```
db2updv10 -d dbname
```

where *dbname* represents the name of the database.

3. Optional: Bind the bind files. Binding of the database utilities and the CLI bind files occurs automatically. However, if an error occurs, you can manually bind the database utilities and the CLI bind files. Refer to “Binding bind files after installing fix packs” on page 484.
4. Optional: Recompile applications.
To take advantage of any changes to the files linked to in the application, recompiling applications is recommended.
5. Optional: If you have installed DB2 Text Search, you must run the **db2iupdt** command with the **/j "TEXT_SEARCH"** option to have it configured.

Results

The fix pack installation and configuration is complete.

Post-installation tasks for fix packs (Linux and UNIX)

As part of a fix pack installation, binding of the database utilities (**IMPORT**, **EXPORT**, **REORG**, the Command Line Processor) and the CLI bind files, DB2 instances are updated automatically.

However, if an error occurs, you can manually bind the database utilities and the CLI bind files and update the DB2 instances. Depending on your database products and the fix pack installation method used, you might need to update the DB2 instances, restart the DB2 instances, restart the DB2 Administration Server and start the **djxlink** command.

Procedure

Perform the following actions:

1. If you have InfoSphere Federation Server installed, run the **djxlink** command.
Perform the following tasks after installing the fix pack and before running **db2iupdt**:
 - a. Log on as root.
 - b. Remove or rename the file **djxlink.out**, which is located in the **DB2DIR/lib** directory, where **DB2DIR** is the DB2 installation directory.
 - c. Ensure that all of the appropriate variables are set, either in your current environment or in the **db2dj.ini** file. For example, if you are using a federated server to connect to an Oracle data source, set the environment variable **ORACLE_HOME** to the Oracle home directory.
 - d. Run the command:

```
djxlink
```
2. Update instances to use the new DB2 database level.
All existing instances in the DB2 copy must be updated after a fix pack is installed. By default, the **installFixPack** command updates the DB2 instances automatically. However, if an error occurs, you can update instances manually.
Perform the following steps:
 - a. Log on as root.

- b. Determine which instances are associated with the DB2 copy by issuing the command:

```
DB2DIR/instance/db2ilist
```

where *DB2DIR* represents the location where the DB2 copy is installed.

- c. If you made any changes to the `db2profile` or `db2cshrc` scripts, either back up the scripts or copy the changes into the `userprofile` and `usercshrc` scripts, respectively.

This action is required because the **db2iupdt** command overwrites the `db2profile` and `db2cshrc` scripts. It does not overwrite the `userprofile` and `usercshrc` scripts.

- d. For each instance, issue the command:

```
DB2DIR/instance/db2iupdt iname
```

where *iname* represents the instance name and *DB2DIR* represents the location where the DB2 copy is installed.

- e. If the DB2 Administration Server (DAS) belongs to the DB2 copy where you installed the fix pack, issue the command:

```
DB2DIR/instance/dasupdt
```

where *DB2DIR* is the location where the DB2 copy is installed. If this DB2 copy is now running at a more recent fix pack level than all of the other DB2 copies, consider updating the DAS to belong to this DB2 copy.

3. Restart the instances and the DAS.

This step is required if you installed a fix pack to update an existing installation. If you installed the fix pack to create a new installation, this step is not required.

To restart an instance:

- a. Log in as the instance owner.
- b. Issue the command **db2start**.

Repeat for each instance.

To restart the DB2 administration server, log in as the DAS owner and run the **db2admin start** command.

4. Optional: If you issued the **db2iauto** command to prevent instances from auto-starting before installing the fix pack, enable auto-start for the instances again. Issue the following command while logged on as root:

```
DB2DIR/instance/db2iauto -on iname
```

where *DB2DIR* is the location where the DB2 copy is installed and *iname* represents the instance owner name. The command must be performed once for each instance that you altered with the **db2iauto** command before you installed the fix pack.

5. Optional: Bind the bind files. Binding of the database utilities and the CLI bind files occurs automatically. However, if an error occurs, you can manually bind the database utilities and the CLI bind files. Refer to “Binding bind files after installing fix packs” on page 484.
6. Optional: Recompile applications.

To take advantage of any changes to the files linked to in the application, recompiling applications is recommended.

Results

Once you have completed these tasks, the fix pack installation and configuration is complete.

Binding bind files after installing fix packs

As part of a fix pack installation on the server, binding of the database utilities (**IMPORT**, **EXPORT**, **REORG**, the Command Line Processor) and the CLI bind files occurs automatically.

However, if you install a fix pack on the client or an error occurs, you can manually bind the database utilities and the CLI bind files. Different subsets of bind files must be bound for DB2 Database for Linux, UNIX, and Windows and host or System i database servers.

Before you begin

Ensure that you have the necessary authority to perform the **BIND** command.

About this task

Note: The IBM Data Server Runtime Client cannot be used to bind the database utilities and CLI bind files. Perform the **BIND** commands from an IBM Data Server Client (or other DB2 database product) that is running on the same operating system and the same DB2 version and fix pack level as the Data Server Runtime Client.

Note: To ensure that not all users have access to databases created with RESTRICTIVE mode, do not grant privileges to PUBLIC for a database created with the RESTRICTIVE mode.

Procedure

To bind the bind files:

1. If you installed the fix pack on DB2 database products that have existing databases, perform the following commands once for each database:

```
db2 terminate
db2 CONNECT TO dbname user USERID using PASSWORD
db2 BIND path\db2schema.bnd BLOCKING ALL GRANT PUBLIC SQLERROR CONTINUE
db2 BIND path\@db2ubind.lst BLOCKING ALL GRANT PUBLIC ACTION ADD
db2 BIND path\@db2cli.lst BLOCKING ALL GRANT PUBLIC ACTION ADD
db2 terminate
```

where *dbname* represents the name of a database to which the files should be bound, and where *path* is the full path name of the directory where the bind files are located, such as *INSTHOME*\sqllib\bnd where *INSTHOME* represents the home directory of the DB2 instance. db2ubind.lst and db2cli.lst contain lists of required bind files used by DB2 database products. Packages that are already bound will return an SQL0719N error. This is expected.

2. Optional: If you installed the fix pack on DB2 database products that have existing databases, rebind the packages by running the **REBIND** or **db2rbind** command.

After you install a fix pack, some packages are marked as invalid. Packages marked as invalid are implicitly rebound the first time an application uses

them. To eliminate this overhead and to ensure that the rebind is successful, manually rebind all packages. For example, issue the **db2rbind** command:

```
db2rbind dbname -l logfile all
```

where *dbname* represents the name of a database whose packages are to be revalidated, and where *logfile* is the name of the file to be used for recording errors during the package revalidation procedure.

3. If you installed the fix pack on DB2 database products that have existing spatial-enabled databases, perform the following commands once for each database:

```
db2 terminate
db2 CONNECT TO dbname
db2 BIND path\BND\@db2gse.1st
db2 terminate
```

where *dbname* represents the name of a database to which the files should be bound, and where *path* is the full path name of the directory where the bind files are located, such as *INSTHOME\sql11ib\bnd* where *INSTHOME* represents the home directory of the DB2 instance. *db2gse.1st* contains the names of the bind files for the stored procedures that DB2 Spatial Extender provides.

4. If you connect to DB2 databases on host or System i servers, perform the following actions:

- For DB2 databases on z/OS® or OS/390®:

```
db2 terminate
db2 CONNECT TO dbname user USERID using PASSWORD
db2 BIND path\@ddcsmv.1st BLOCKING ALL SQLERROR CONTINUE GRANT PUBLIC ACTION ADD
db2 terminate
```

- For DB2 databases on VM:

```
db2 terminate
db2 CONNECT TO dbname user USERID using PASSWORD
db2 BIND path\@ddcsvm.1st BLOCKING ALL SQLERROR CONTINUE GRANT PUBLIC ACTION ADD
db2 terminate
```

- For DB2 databases on VSE:

```
db2 terminate
db2 CONNECT TO dbname user USERID using PASSWORD
db2 BIND path\@ddcsvse.1st BLOCKING ALL SQLERROR CONTINUE GRANT PUBLIC ACTION ADD
db2 terminate
```

- For DB2 databases on System i:

```
db2 terminate
db2 CONNECT TO dbname user USERID using PASSWORD
db2 BIND path\@ddcs400.1st BLOCKING ALL SQLERROR CONTINUE GRANT PUBLIC ACTION ADD
db2 terminate
```

where *dbname* represents the name of a host or System i database to which the files should be bound, and where *path* is the full path name of the directory where the bind files are located, such as *INSTHOME\sql11ib\bnd* where *INSTHOME* represents the home directory of the DB2 instance.

5. If you connect to databases that are running on different operating systems (Linux, UNIX or Windows) or at different DB2 versions or service levels, bind the database utilities and CLI bind files against those databases.

Note:

- The actions required are the same irrespective of whether you connect to a database on another DB2 database system or in another DB2 copy on the same machine.
- If you have installed the fix pack in multiple locations, perform the actions once from each unique combination of operating system and DB2 version or service level.

Perform the following actions:

```
db2 terminate
db2 CONNECT TO dbname user USERID using PASSWORD
db2 BIND path\@db2ubind.1st BLOCKING ALL GRANT PUBLIC ACTION ADD
db2 BIND path\@db2cli.1st BLOCKING ALL GRANT PUBLIC ACTION ADD
db2 terminate
```

where *dbname* represents the name of a database to which the files should be bound, and where *path* is the full path name of the directory where the bind files are located, such as *INSTHOME*\sql1lib\bnd where *INSTHOME* represents the home directory of the instance where you are issuing the commands. *db2ubind.1st* and *db2cli.1st* contain lists of required bind files used by DB2 database products. Packages that are already bound will return an SQL0719N error. This is expected.

Binding federated databases

If you have existing federated databases, you must bind the bind files *db2dsproc.bnd* and *db2stats.bnd* after you install a DB2 fix pack. To bind the bind files, you must have one of the following authorities:

- DBADM authority
- ALTERIN privilege on the schema
- BIND privilege on the package

To bind the bind files *db2dsproc.bnd* and *db2stats.bnd*, connect to the database and run the **BIND** command. For example:

```
db2 CONNECT TO dbname user USERID using PASSWORD
db2 bind path/db2dsproc.bnd blocking all grant public
db2 bind path/db2stats.bnd blocking all grant public
db2 terminate
```

where *dbname* represents the name of the federated database, and *path* represents the full path name of the directory where the bind files are located, such as *\$HOME*/sql1lib/bnd where *\$HOME* represents the DB2 instance home directory.

Chapter 55. Uninstalling fix packs

After installing a fix pack, you can revert to a previous fix pack or the GA level of the DB2 database product.

Before you begin

On Linux and UNIX operating systems, you can go back to an earlier fix pack or the GA level.

On Windows operating systems, you cannot go back to a previous fix pack or the GA level unless you uninstall the current fix pack, then install the previous level.

Procedure

1. On Linux or UNIX operating systems, to uninstall a fix pack, use the **installFixPack** command with the force option (**-f**) to bypass the level check. The command must be run from the lower level fix pack or GA image. For example:

```
./installFixPack -f level -b DB2DIR
```

where

- *DB2DIR* is the location of the DB2 database product that you want to force to a lower level fix pack or GA image. For example:

```
./installFixPack -f level -b /opt/ibm/db2/V10.1
```

2. On Windows operating system, to uninstall a fix pack, use the Add/Remove Programs window, accessible through the Windows Control Panel. Refer to your operating system's help for more information about removing software products from your Windows operating system.
3. Associate the instance with another DB2 copy that is at a lower level than the original DB2 copy where the instance was running. You can update all instances by issuing **db2iupdt -f level** from the new directory.

What to do next

Note: Before you uninstall the current fix pack, you remove or uninstall any functionality that does not apply to the earlier fix pack or the GA level.

Part 10. Uninstalling DB2 database products

Removing a DB2 database product from your operating system.

Only completely uninstall your DB2 database product from your Windows operating system if you no longer require your existing DB2 instances and databases.

The uninstall process on your Linux and UNIX operating system is different for root and non-root DB2 database products

Chapter 56. Uninstalling your DB2 database product (Linux and UNIX)

This task provides steps for removing a DB2 database product from your Linux or UNIX operating system.

About this task

This task is not required to install a new version of a DB2 database product. Each version of a DB2 database product on Linux or UNIX has a different installation path and can therefore coexist on the same computer.

Note: This task applies to DB2 database products that were installed with root user authority. A separate topic explains how to uninstall DB2 database products that were installed as a non-root user.

Procedure

To remove your DB2 database product:

1. Optional: Drop all databases. You can drop databases using the **DROP DATABASE** command. Database files remain intact on your file systems when you drop an instance without dropping databases first.
2. Stop the DB2 Administration Server. Refer to the *Installing DB2 Servers* manual.
3. Remove the DB2 Administration Server, or run the **dasupdt** command to update the DB2 Administration Server to another installation path. To remove the DB2 Administration Server, refer to the *Installing DB2 Servers* manual.
4. Stop all DB2 instances. Refer to the *Installing DB2 Servers* manual.
5. Remove the DB2 instances, or run the **db2iupdt** command to update the instances to another installation path. To remove the DB2 instances, refer to the *Installing DB2 Servers* manual.
6. Remove the DB2 database products. Refer to the *Installing DB2 Servers* manual.

Stopping the DB2 administration server (Linux and UNIX)

You must stop the DB2 administration server (DAS) before you remove your DB2 product.

About this task

Important: The DB2 Administration Server (DAS) has been deprecated in Version 9.7 and might be removed in a future release. The DAS is not supported in DB2 pureScale environments. Use software programs that use the Secure Shell protocol for remote administration. For more information, see “DB2 administration server (DAS) has been deprecated” at .

When uninstalling a DB2 product, you must drop the DAS if you are removing your last DB2 copy. If you have other DB2 copies, it is recommended that you run the **dasupdt** command to associate the DAS with another DB2 copy. If you decide to drop the DAS, you must stop the DAS first.

Note: This task does not apply to non-root installations of DB2 products.

Procedure

To stop the DB2 administration server:

1. Log in as the DB2 administration server owner.
2. Stop the DB2 administration server by entering the **db2admin stop** command.

Removing the DB2 administration server (Linux and UNIX)

If you are removing your last DB2 copy, you must remove the DB2 administration server (DAS) before you remove your DB2 database product.

About this task

Important: The DB2 Administration Server (DAS) has been deprecated in Version 9.7 and might be removed in a future release. The DAS is not supported in DB2 pureScale environments. Use software programs that use the Secure Shell protocol for remote administration. For more information, see “ DB2 administration server (DAS) has been deprecated” at .

If you are removing a DB2 copy, but have other DB2 copies, run the **dasupdt** command from the DB2 copy that you want to have the DB2 DAS associated with.

Restrictions

This task applies only to DB2 database products that were installed with root user authority.

Procedure

To remove the DAS:

1. Log in as a user with root user authority.
2. Stop the DAS. For example:

```
db2admin stop
```
3. Remove the DAS. Enter the following command:

```
DB2DIR/instance/dasdrop
```

where *DB2DIR* is the location you specified during the DB2 database product installation. The default installation path for UNIX is `/opt/IBM/db2/V10.1`. The default installation path for Linux is `/opt/ibm/db2/V10.1`.

Stopping root DB2 instances (Linux and UNIX)

You must stop all DB2 instances associated with the DB2 copy you are uninstalling. Instances associated with other DB2 copies might not be affected by uninstalling the current copy.

About this task

Note: This task applies to DB2 database products that were installed with root user authority.

Procedure

To stop a DB2 instance:

1. Log in as a user with root user authority.
2. Obtain a list of the names of all DB2 instances associated with your current DB2 copy by entering the following command:

```
DB2DIR/bin/db2ilist
```

where *DB2DIR* is the location you specified during the DB2 database product installation. The default installation path for UNIX is */opt/IBM/db2/V10.1*. The default installation path for Linux is */opt/ibm/db2/V10.1*.

3. Run the script if it is not included in *.profile*.

```
. INSTHOME/sql1lib/db2profile      (bash, Bourne, or Korn shells)
source INSTHOME/sql1lib/db2cshrc  (C shell)
```

where *INSTHOME* is the home directory of the instance.

4. It is recommended that you save the following files:
 - The database manager configuration file, *\$HOME/sql1lib/db2system*
 - The node configuration file, *\$HOME/sql1lib/db2nodes.cfg*
 - User defined functions or fenced stored procedure applications in *\$HOME/sql1lib/function*
5. Stop the DB2 database manager by entering the **db2stop force** command.
6. Confirm that the instance is stopped by entering the **db2 terminate** command.
7. Repeat these steps for each instance.

Removing DB2 instances (Linux and UNIX)

This task explains how to remove some or all of the root instances on your system. Remove DB2 instances only if you are not planning to use your DB2 database products, or if you do not want to upgrade existing instances to a later version of the DB2 database product.

About this task

If you are removing your last DB2 Version 9 copy, you can remove the DB2 instances before you remove your DB2 database product. If you removing a DB2 Version 9 copy, but have other DB2 Version 9 copies, you can run the **db2iupdt** command from the DB2 copy that you would like to have the DB2 instances associated with.

Once an instance is removed, you can use the DB2 databases owned by the instance if you catalog them under another instance of the same release. Even though you remove the instance, the databases are still intact and can be reused unless the database files are expressly deleted.

Upgrading requires that both the new and old DB2 database versions are still installed. You cannot upgrade an instance if the DB2 copy it is associated with has been removed.

Restrictions

This task does not apply to non-root installations. To remove a non-root instance, you must uninstall your DB2 database product.

Procedure

To remove an instance:

1. Log in as a user with root user authority.
2. Optional: If you are certain you no longer need the data in the associated databases, you can remove the database files from the systems or drop the databases before dropping the instance.
3. Remove the instance by entering the following command:

```
DB2DIR/instance/db2idrop InstName
```

where *DB2DIR* is the location you specified during the DB2 database product installation. The default installation path for UNIX is `/opt/IBM/db2/V10.1`. The default installation path for Linux is `/opt/ibm/db2/V10.1`.

The **db2idrop** command removes the instance entry from the list of instances and removes the *INSTHOME*/sql1ib directory, where *INSTHOME* is the home directory of the instance and where *InstName* is the login name of the instance. If you are storing any files in `/sql1ib` directory, these files will be removed by this action. If you still need these files, you must make a copy of them before dropping the instance.

4. Optional: As a user with root user authority, remove the instance owner's user ID and group (if used only for that instance). Do not remove these if you are planning to re-create the instance.

Note: This step is optional since the instance owner and the instance owner group might be used for other purposes.

Removing DB2 database products using the **db2_deinstall** and **doce_deinstall** commands (Linux and UNIX)

This task provides steps for removing DB2 database products or DB2 database components using the **db2_deinstall** and **doce_deinstall** commands.

Before you begin

Before you remove DB2 database products from your system, ensure that you have performed all steps outlined in Chapter 56, “Uninstalling your DB2 database product (Linux and UNIX),” on page 491

About this task

This task applies to DB2 database products that were installed with root user authority.

The **db2_deinstall** command removes DB2 database products from your system.

The **doce_deinstall** command removes the *DB2 Information Center* that is in the same install path as the **doce_deinstall** tool.

Restrictions

- You cannot remove DB2 database products using a native operating system utility, such as **rpm** or **SMIT**.
- The **doce_deinstall** command is available only on Linux operating systems (Linux x32 and x64).

Procedure

To remove DB2 database products, features or the *DB2 Information Center* from a specific path:

1. Log in with root user authority.
2. Access the path where the DB2 database products are located.
3. Run one of the following commands:
 - To remove a feature from an installed DB2 database product in the current location, run the **db2_deinstall -F** command from the *DB2DIR/install* directory.
 - To remove all installed DB2 database products in the current location, run the **db2_deinstall -a** command from the *DB2DIR/install* directory.
 - To remove a DB2 database product using a response file, run the **db2_deinstall -r *response_file*** command from the *DB2DIR/install* directory. You can use a sample response file to uninstall the product. For example, **doce_deinstall -r db2un.rsp**
 - To remove the *DB2 Information Center* in the current location, run the **doce_deinstall -a** from the *DB2DIR/install* directory.
 - To remove the *DB2 Information Center* using a response file, run the **doce_deinstall -r *response_file*** from the *DB2DIR/install* directory. You can use the sample response file to uninstall the Information Center. For example, **doce_deinstall -r doceun.rsp**

where *DB2DIR* is the location that you specified when you installed your DB2 database product.

Chapter 57. Uninstalling non-root DB2 database products (Linux and UNIX)

This task provides steps for removing non-root DB2 database products from your Linux or UNIX operating system.

About this task

Note: This task applies to DB2 database products that were installed without root user authority. To uninstall DB2 database products that were installed with root user authority, see Chapter 56, “Uninstalling your DB2 database product (Linux and UNIX),” on page 491

Non-root installations of DB2 database products are uninstalled in the same manner as root installations. However, there are important differences which are detailed in the substeps.

Procedure

To remove your DB2 database product:

1. Stop your non-root instance.
2. Remove your DB2 database product.

Stopping non-root instances (Linux and UNIX)

You must stop your non-root instance before uninstalling your DB2 database product.

About this task

Note: This task applies to DB2 database products that were installed with non-root authority.

Procedure

To stop a DB2 instance:

1. Log in as the non-root instance owner.
2. Run the start up script if it is not included in `.profile`.

```
. $HOME/sqlllib/db2profile      (bash, Bourne, or Korn shells)
source $HOME/sqlllib/db2cshrc   (C shell)
```

where `$HOME` is your home directory.

3. You might want to save any of the following files:
 - The database manager configuration file, `db2system`
 - The configuration file used to enable root features before running **db2rfe**.
 - User defined functions or fenced stored procedure applications in `$HOME/sqlllib/function`.
4. Stop the DB2 database manager by entering the **db2stop force** command.
5. Confirm that the instance is stopped by entering the **db2 terminate** command.

Removing non-root DB2 database products using `db2_deinstall` (Linux and UNIX)

This task provides steps for removing non-root DB2 database products or components using the `db2_deinstall` command.

Before you begin

You must stop the non-root instance before running the `db2_deinstall` command.

About this task

- This task applies to DB2 database products that were installed without root user authority. A separate task exists for uninstalling DB2 database products that were installed with root user authority.
- As with root users, non-root users can use the `db2_deinstall` command to uninstall DB2 database products. The `db2_deinstall` command for non-root installations has the same options as root installations, and has an extra option: `-f sqllib`.
- It is important to note that running `db2_deinstall` as a non-root user uninstalls the DB2 database product *and* drops the non-root instance. This is different than root installations, where running `db2_deinstall` only uninstalls the DB2 database program files.
- You cannot remove DB2 database products using a native operating system utility, such as `rpm` or `SMIT`.

Procedure

To uninstall a DB2 database product that was installed by a non-root user:

1. Log in with the user ID that was used to install the DB2 database product.
2. Navigate to the `$HOME/sqllib/install` directory, where `$HOME` is your home directory.
3. Run the `db2_deinstall` command.

Note:

- If you run the `db2_deinstall` command with the `-a` option, the DB2 database program files are removed, but any configuration files are left behind in a backup directory called `sqllib_bk`.
- If you run the `db2_deinstall` command with the `-a -f sqllib` option, the entire `sqllib` subdirectory in your home directory will be removed. If you have any files in `sqllib` you want to keep, be sure to copy them elsewhere before running `db2_deinstall -a -f sqllib`.
- As with root installations, running the `db2_deinstall` command with the `-F` option against a non-root installation allows the non-root user to remove specific DB2 features.

Chapter 58. Uninstalling your DB2 database product (Windows)

This task provides steps for completely removing your DB2 database product from your Windows operating system. Only perform this task if you no longer require your existing DB2 instances and databases.

About this task

If you are uninstalling the default DB2 copy, and you have other DB2 copies on your system, use the **db2swtch** command to choose a new default copy before you proceed with the uninstallation. Also, if your DB2 Administration Server (DAS) is running under the copy being removed, move your DAS to a copy that is not being removed. Otherwise, re-create the DAS using the **db2admin create** command after the uninstall, and you reconfigure the DAS for some function to work.

Procedure

To remove your DB2 database product from Windows:

1. Optional: Drop all databases using the **drop database** command. Be sure that you no longer need these databases. If you drop your databases, all of your data will be gone.
2. Stop all DB2 processes and services. This can be done through the Windows Services panel or by issuing the **db2stop** command. If DB2 services and processes are not stopped before attempting to remove your DB2 database product, you will receive a warning containing a list of processes and services that are holding DB2 DLLs in memory. If you will use Add/Remove Programs to remove your DB2 database product, this step is optional.
3. You have two options for removing your DB2 database product:
 - Add/Remove Programs
Accessible through the Windows Control Panel, use the Add/Remove Programs window to remove your DB2 database product. Refer to your operating system's help for more information about removing software products from your Windows operating system.
 - **db2unins** command
You can run the **db2unins** command from the *DB2DIR\bin* directory to remove your DB2 database products, features, or languages. Using this command, you can uninstall multiple DB2 database products at the same time using the **/p** parameter. You can use a response file to uninstall DB2 database products, features, or languages using **/u** parameter.

What to do next

Unfortunately, your DB2 database product cannot always be removed by using the **Control Panel > Add/Remove Programs** facility or using the **db2unins /p** command or the **db2unins /u** command. The following uninstallation option must ONLY be attempted if the previous method fails.

To forcefully remove all DB2 copies from your Windows system, run the **db2unins /f** command. This command will perform a brute force uninstallation of ALL DB2 copies on the system. Everything except user data, such as DB2 databases, will be

forcefully deleted. Before running this command with the **/f** parameter, see the **db2unins** command for details.

Chapter 59. Uninstalling DB2 copies and IBM database client interface copies

This task will guide you through the process of uninstalling DB2 copies and IBM database client interface copies. Please note that there are notes for both Windows platforms and Linux and UNIX platforms.

About this task

On Linux and UNIX operating systems

Use the **db2_deinstall** command from the DB2 copy that you are using. The **db2_deinstall** command uninstalls installed DB2 products or features that are in the same install path as the **db2_deinstall** tool.

The **db2_deinstall** command is also available on the DB2 product DVD. The product version you are uninstalling must match the product version on the DB2 DVD. If you do not provide the **-b** parameter when you run the **db2_deinstall** command from the DB2 DVD, you will be prompted for the installation path.

Use the **db2ls** command to see the list of installed DB2 products and features. If one or more instances are currently associated with a DB2 copy, that DB2 copy cannot be uninstalled.

On Windows operating systems

To uninstall DB2 copies on Windows operating systems, use one of the following methods:

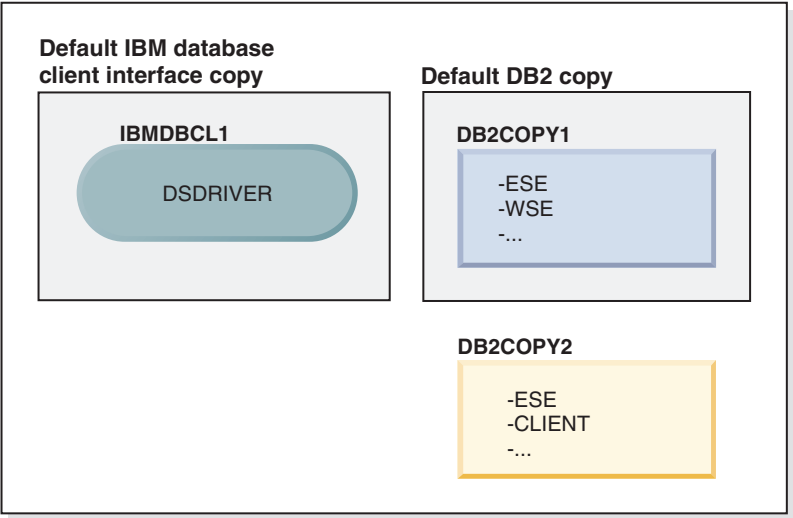
- Use the Windows Add/Remove Control Panel Applet.
- Run the **db2unins** command from the installed DB2 copy directory

Note:



- You can uninstall DB2 even when there are instances associated with DB2 copies. If you do this, the instance information will be removed with the DB2 uninstall. Therefore, take extra care when managing, recovering, and uninstalling instances.
- If multiple Version 9 copies are installed, you cannot remove the default DB2 copy. If you want to remove the default DB2 copy, switch the default DB2 copy to one of the other DB2 copies before uninstalling. For more information about switching the default DB2 copy, see the **db2swtch** command.

Uninstalling DB2 copies when both those copies and IBM Data Server Driver copies are on your machine

DB2 Version 9.5

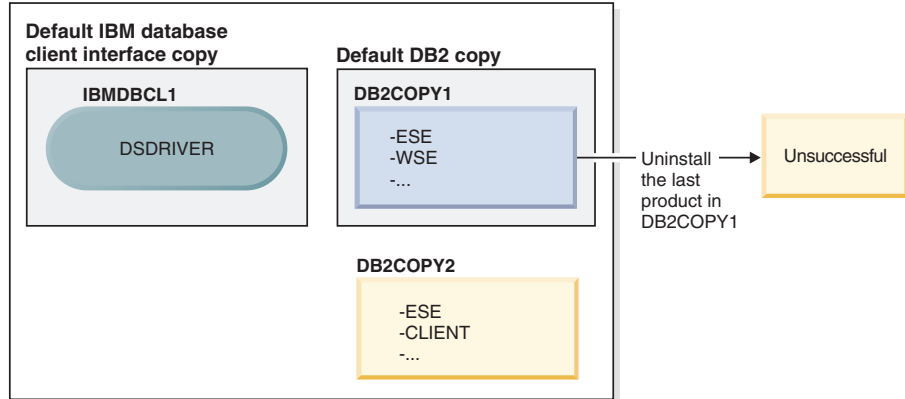


Legend



	Default DB2 copy
	Default IBM database client interface copy
DSDRIVER = IBM Data Server Driver for ODBC, CLI, and .NET	

In the scenario presented here, IBMDBCL1 is the default IBM database client interface copy, DB2COPY1 is the default DB2 copy, and there is another DB2 copy (DB2COPY2).

DB2 Version 9.5




Legend

	Default DB2 copy
	Default IBM database client interface copy
DSDRIVER = IBM Data Server Driver for ODBC, CLI, and .NET	

You decide to uninstall all of the DB2 products that are part of DB2COPY1. When you attempt to uninstall the last DB2 product within DB2COPY1, because it is the default DB2 copy, the uninstall request will be unsuccessful.

Before uninstalling the last of the remaining DB2 products within the default DB2 copy, and if there is another DB2 copy on your system, you must switch the default from one DB2 copy to the other.

DB2 Version 9.5



Legend



Default DB2 copy



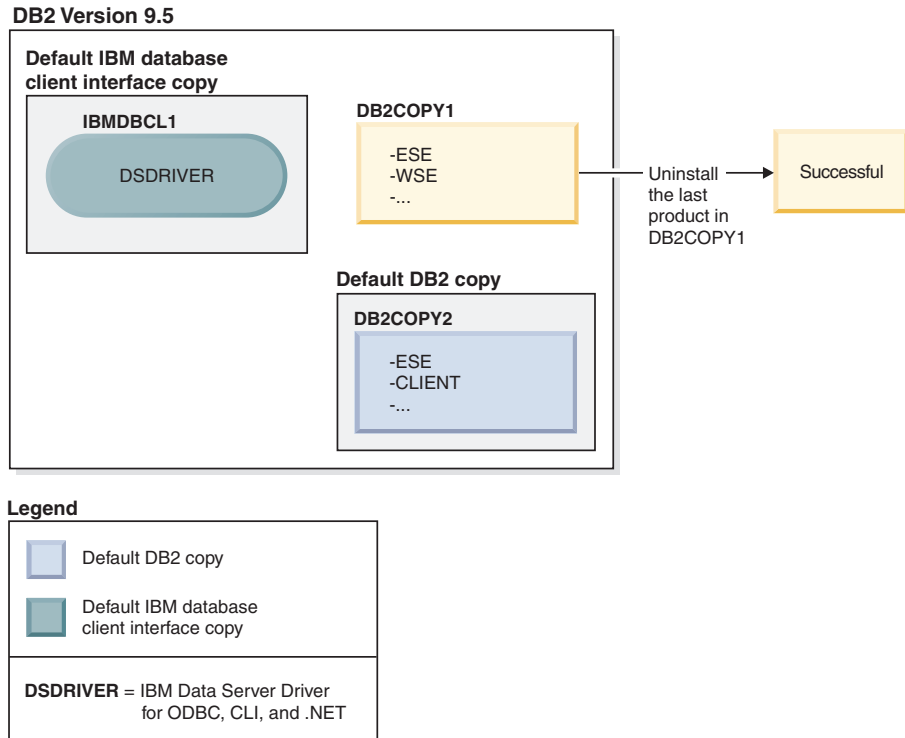
Default IBM database client interface copy

DSDRIVER = IBM Data Server Driver for ODBC, CLI, and .NET

To change the default, use the `db2swtch` command with no arguments (Windows only) to start the Default DB2 and IBM Database Client Interface Selection wizard. The wizard will show all possible candidates when choosing a new default.

In this case, you can choose DB2COPY2 as the new default DB2 copy.

After you have made DB2COPY2 the default DB2 copy, you can request the uninstall of the last product in DB2COPY1.

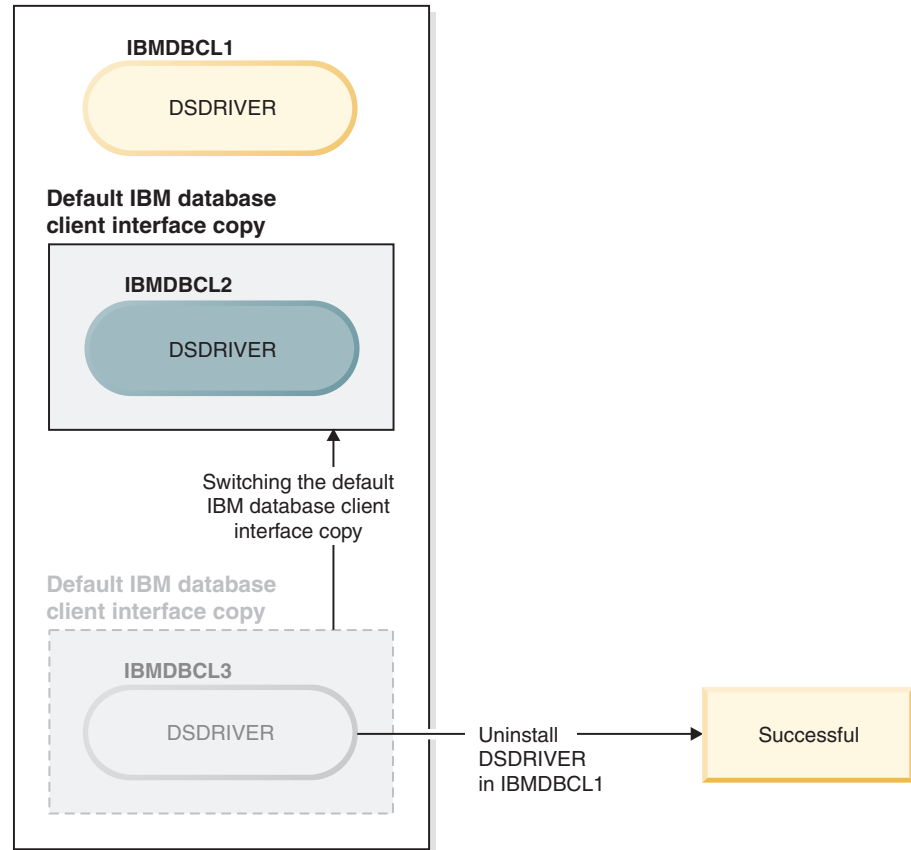


Since DB2COPY1 is no longer the default DB2 copy, the uninstall request is successful.


Uninstalling IBM Data Server Driver copies when only those copies are on your machine

Over time you can have several DSDRIVERS installed. Only one IBM database client interface copy is the default. At some point you might decide to uninstall the DSDRIVER that is the default IBM database client interface copy.

DB2 Version 9.5



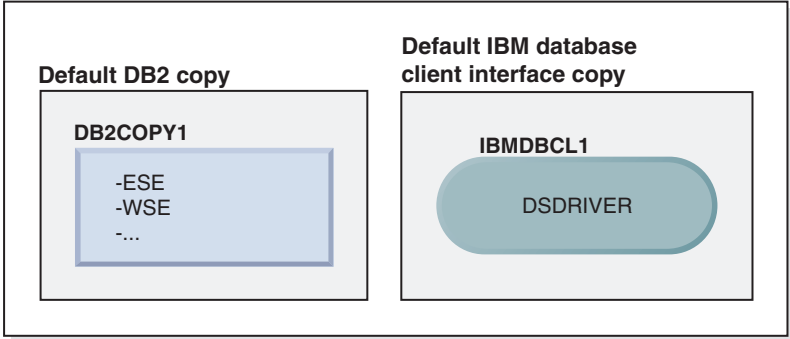
Legend

	Default IBM database client interface copy
DSDRIVER = IBM Data Server Driver for ODBC, CLI, and .NET	


When you uninstall the default IBM database client interface copy, the database manager will choose between any of the remaining DSDRIVERS and pick one of them to be the new default IBM database client interface copy. Unless you switch the default before requesting the uninstall, you will not control which DSDRIVER will become the new default. (If there is only one other DSDRIVER besides the original default, then you will know which DSDRIVER the database manager will choose. If there are several DSDRIVERS installed besides the original default, then you will not know which DSDRIVER the database manager will choose.)

Uninstalling IBM Data Server Driver copies when both those copies and DB2 copies are on your machine

DB2 Version 9.5



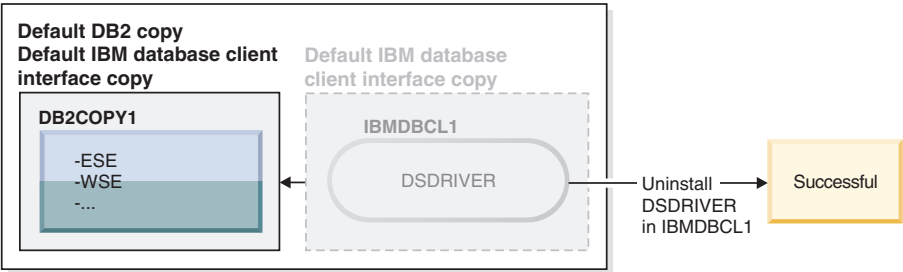
Legend




In the scenario presented here, DB2COPY1 is the default DB2 copy and IBMDBCL1 is the default IBM database client interface copy.

You decide to uninstall the DSDRIVER in IBMDBCL1.

DB2 Version 9.5



Legend



As part of the uninstall request, the default IBM database client interface copy is switched by the database manager automatically so that DB2COPY1 becomes both the default DB2 copy and the default IBM database client interface copy. (This is also be true if there is more than one DB2 copy on the machine.)

Chapter 60. Uninstalling a DB2 database product, feature, or language using a response file (Linux and UNIX)

To silently uninstall DB2 database products, features, or languages in a DB2 copy, use the **db2_deinstall** command with the **-r** option.

About this task

You can also use a response file to uninstall the *DB2 Information Center*.

If you have DB2 database products installed in various installation paths, you must run this command separately from each installation path. Sample response file, `db2un.rsp`, is in `DB2DIR/install`, where `DB2DIR` specifies the path where the DB2 database product was installed.

Procedure

To perform an uninstallation:

1. Customize the `db2un.rsp` response file. To activate an item in the response file, remove the asterisk (*) to the left of the keyword. Then, replace the current setting to the right of the value with the new setting. The possible settings are listed to the right of the equal sign.
2. Run the **db2_deinstall** command. For example, **db2_deinstall -r db2un.rsp**.
3. Check the messages in the log file when the uninstallation finishes. The log file is located in:
 - For root installations: `/tmp/db2_deinstall.log.process-id`
 - For non-root installations: `/tmp/db2_deinstall_user-id.log`

Chapter 61. Uninstalling a DB2 product, feature, or language using a response file (Windows)

These instructions guide you through the process of uninstalling a DB2 product, feature, or language using a response file on a Windows platform.

Before you begin

Before you begin the uninstallation, ensure that you have all the required user accounts to perform the uninstallation.

About this task

To silently uninstall DB2 products, features, or languages in a DB2 copy, use the **db2unins** command with the **-u** parameter.

If you have DB2 products installed in various installation paths, you must run this command separately from each installation path. Sample response file, **db2un.rsp**, is in *DB2DIR\install*, where *DB2DIR* specifies the path where the DB2 product was installed.

If you have a clustered environment, before uninstalling your DB2 product you must run the **db2mscs** command, with the **-u** parameter, from the same server that originally ran the **db2mscs** command to create the failover infrastructure. For details, see the **db2mscs** command.

Procedure

To perform an uninstallation:

1. Customize the **db2un.rsp** response file. To activate an item in the response file, remove the asterisk (*) to the left of the keyword. Then, replace the current setting to the right of the value with the new setting. The possible settings are listed to the right of the equal sign.
2. Run the **db2unins** command. For example, **db2unins -u c:\db2un.rsp**. The **db2unins** command is found under the *SQLLIB\BIN* directory.
3. Check the messages in the log file when the uninstallation finishes. The default log file location is *My Documents\DB2LOG\db2un_timestamp.log*.

Chapter 62. Uninstalling fix packs

After installing a fix pack, you can revert to a previous fix pack or the GA level of the DB2 database product.

Before you begin

On Linux and UNIX operating systems, you can go back to an earlier fix pack or the GA level.

On Windows operating systems, you cannot go back to a previous fix pack or the GA level unless you uninstall the current fix pack, then install the previous level.

Procedure

1. On Linux or UNIX operating systems, to uninstall a fix pack, use the **installFixPack** command with the force option (**-f**) to bypass the level check. The command must be run from the lower level fix pack or GA image. For example:

```
./installFixPack -f level -b DB2DIR
```

where

- *DB2DIR* is the location of the DB2 database product that you want to force to a lower level fix pack or GA image. For example:

```
./installFixPack -f level -b /opt/ibm/db2/V10.1
```

2. On Windows operating system, to uninstall a fix pack, use the Add/Remove Programs window, accessible through the Windows Control Panel. Refer to your operating system's help for more information about removing software products from your Windows operating system.
3. Associate the instance with another DB2 copy that is at a lower level than the original DB2 copy where the instance was running. You can update all instances by issuing **db2iupdt -f level** from the new directory.

What to do next

Note: Before you uninstall the current fix pack, you remove or uninstall any functionality that does not apply to the earlier fix pack or the GA level.

Part 11. Installing and upgrading SA MP with the DB2 installer

IBM Tivoli System Automation for Multiplatforms (SA MP) is integrated with IBM DB2 server as part of the DB2 High Availability Feature on AIX, Linux, and Solaris operating systems. You can install, upgrade, or uninstall SA MP using either the DB2 installer or the **installSAM** and **uninstallSAM** scripts that are included in the DB2 server install media. On Windows operating systems, the SA MP is bundled as part of the DB2 High Availability Feature, but it is not integrated with the DB2 installer.

Before you begin

- To install and use SA MP, your system configuration and intended use of SA MP must meet the terms of the license that comes with the SA MP that is integrated with DB2 server.

For information about the license details of SA MP that is integrated with DB2 server, see Chapter 69, “License terms for using IBM Tivoli System Automation for Multiplatforms (SA MP) integrated with IBM DB2 server,” on page 537.

- To install or upgrade SA MP your system architecture must be supported by the SA MP that is integrated with DB2 server.

For more information about SA MP supported software and hardware, see Chapter 70, “Supported software and hardware for IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 539.

- You must have root authority to install SA MP.

If you perform a non-root installation of DB2 server, you can install SA MP from the IBM DB2 installation media separately. When you install SA MP separately, you still must have root authority.

- SA MP does not support AIX system workload partitions (WPARs), Solaris 9 platform, Solaris x64 AMD64, or non-global zones on Solaris. For a list of supported platforms and system configurations, see <http://publib.boulder.ibm.com/tividd/td/IBMTivoliSystemAutomationforMultiplatforms3.1.html>.
- If a IBM PowerHA SystemMirror for AIX cluster is running, you cannot perform a SA MP installation or upgrade because SA MP bundles Reliable Scalable Cluster Technology (RSCT) filesets that are dependent on PowerHA SystemMirror. To skip the SA MP installation use the **db2_install** command or the **installFixPack** command.

For information about installing or upgrading SA MP using a PowerHA SystemMirror cluster, see the white paper entitled “Upgrade guide for DB2 Servers in HACMP Environments”, which is available from the IBM Support and downloads website (<http://www.ibm.com/support/docview.wss?uid=swg21045033>).

Chapter 63. IBM Tivoli System Automation for Multiplatforms (SA MP) base component

IBM Tivoli System Automation for Multiplatforms (SA MP) provides high availability and disaster recovery capabilities for AIX, Linux, Solaris SPARC, and Windows.

SA MP is integrated with DB2 Enterprise Server Edition, DB2 Advanced Enterprise Server Edition, DB2 Workgroup Server Edition, DB2 Connect Enterprise Edition and DB2 Connect Application Server Edition on AIX, Linux, and Solaris SPARC operating systems. It is also integrated with Express Edition for use with DB2 Express-C Fixed Term License (FTL) and the DB2 High Availability Feature.

On Windows operating systems, SA MP is bundled with all of these DB2 database products and features, but it is not integrated with the DB2 database product installer.

You can use this copy of SA MP to manage the high availability of your DB2 database system. You cannot use this copy to manage database systems other than DB2 database systems without buying an upgrade for the SA MP license.

SA MP is the default cluster manager in an IBM DB2 server clustered environment on AIX, Linux, and Solaris SPARC operating systems.

For more information about SA MP, see IBM Tivoli System Automation for Multiplatforms (SA MP) publib.boulder.ibm.com/tividd/td/IBMTivoliSystemAutomationforMultiplatforms3.1.html. The list of supported operating systems is also available on the following website: www.ibm.com/software/tivoli/products/sys-auto-linux/platforms.html.

Chapter 64. Installing IBM Tivoli System Automation for Multiplatforms (SA MP)

You can install IBM Tivoli System Automation for Multiplatforms (SA MP) using either the DB2 installer or the **installSAM** installation script that is included in the IBM DB2 server installation media.

Before you begin

Whether you are using the DB2 installer, **installSAM**, or **uninstallSAM**, you must meet the basic prerequisites for installing, upgrading, or uninstalling SA MP. See Part 11, “Installing and upgrading SA MP with the DB2 installer,” on page 515.

If you already have SA MP installed, you can upgrade the installed version of SA MP using the DB2 installer or the **installSAM** installation script. For more information about upgrading SA MP, see Chapter 65, “Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 523.

If a IBM PowerHA SystemMirror for AIX cluster is running, you cannot perform a SA MP installation or upgrade because SA MP bundles Reliable Scalable Cluster Technology (RSCT) filesets that are dependent on PowerHA SystemMirror. To skip the SA MP installation use the **db2_install** command or the **installFixPack** command. For information about installing or upgrading SA MP using a PowerHA SystemMirror cluster, see the white paper entitled Upgrade guide for “DB2 Servers in HACMP Environments”, which is available from the IBM Support and downloads website (<http://www.ibm.com/support/docview.wss?uid=swg21045033>).

Procedure

There are two methods for installing or upgrading SA MP:

- Using the DB2 installer.
- Using the **installSAM** install script that is included in the IBM Data Server install media.

What to do next

Look in the SA MP installation log for diagnostic information about any warnings or errors that the DB2 installer or the **installSAM** installation script returned. For more information about the SA MP installation log, see Chapter 68, “IBM Tivoli System Automation for Multiplatforms (SA MP) install and uninstall logs,” on page 535.

Installing IBM Tivoli System Automation for Multiplatforms (SA MP) using the DB2 installer

You can install IBM Tivoli System Automation for Multiplatforms (SA MP) using the DB2 installer.

Before you begin

Whether you use the DB2 installer or the **installSAM** installation script to install SA MP, you must meet the basic prerequisites for installing SA MP. See: Chapter 64, “Installing IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 519.

About this task

There are three methods for using the DB2 installer to install or uninstall:

- DB2 Setup wizard (install, upgrade, or uninstall)
- Silent installation by using a response file with **db2setup** (install or upgrade) or **db2unins** (for uninstall)
- **db2_install** command (for install), **installFixPack** command (for upgrade), or **db2_deinstall** command (for uninstall)

Before installing SA MP on a given host, the DB2 installer queries your system for the following information:

- Is SA MP on your IBM Data Server installation media?
- Is SA MP already installed?

The DB2 installer calls the **installSAM** installation script to perform some parts of the SA MP installation operation. Instead of using the DB2 installer to install SA MP, you can call **installSAM** directly. For more information about using the **installSAM** installation script to install SA MP, see: “Installing IBM Tivoli System Automation for Multiplatforms (SA MP) using the installSAM install script” on page 521.

You can use the **-l** option with **db2setup**, **db2_install**, or **installFixPack** to specify where the **installSAM** script saves the SA MP installation log. For more information about the SA MP installation log, see: Chapter 68, “IBM Tivoli System Automation for Multiplatforms (SA MP) install and uninstall logs,” on page 535.

Procedure

- To install SA MP using DB2 Setup wizard, run DB2 Setup wizard and follow the instructions in the wizard.

The information about your system that the DB2 installer collects determines which panels appear in the graphical interface of the DB2 Setup wizard during install. For example, if you already have SA MP installed, then the DB2 Setup wizard will not display a panel to install SA MP.

- To install SA MP using a response file, set the response file keyword **INSTALL_TYPE** to "CUSTOM" and then set **COMP** to "TSAMP".

In a response file installation operation, the default behavior of the DB2 installer is to install SA MP. If **INSTALL_TSAMP** is "YES", or if **INSTALL_TSAMP** is commented out or missing from the response file, the DB2 installer installs SA MP.

To prevent the DB2 installer from installing SA MP in a response file installation, set **INSTALL_TSAMP** to "NO".

Important: The **INSTALL_TSAMP** response file keyword has been deprecated in Version 10.1 and might be removed in a future release. Starting in Version 10.1, IBM Tivoli System Automation for Multiplatforms (SA MP) is automatically installed when SA MP is a required component for other components selected

for installation. For more information, see “ IBM Tivoli System Automation for Multiplatforms (SA MP) is now automatically installed” in *What’s New for DB2 Version 10.1*.

- To install SA MP using the **db2_install** command or the **installFixPack** command, you can run the commands without any parameters specific to SA MP.

The default behavior is to install SA MP.

To prevent the SA MP installation, use the **-f NOTSAMP** option.

What to do next

Whether you use the DB2 installer or the **installSAM** installation script to install SA MP, follow the same post-installation steps. For more information about general post-install steps, see: Chapter 64, “Installing IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 519

Installing IBM Tivoli System Automation for Multiplatforms (SA MP) using the installSAM install script

You can install IBM Tivoli System Automation for Multiplatforms (SA MP) using the **installSAM** installation script that is included in the IBM DB2 installation media.

Before you begin

Whether you use the DB2 installer or the **installSAM** installation script to install SA MP, you must meet the basic prerequisites for installing SA MP. See: Chapter 64, “Installing IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 519.

Procedure

Run the **installSAM** installation script

The **installSAM** script is located on the IBM DB2 install media at the following location:

```
db2/platform/tsamp
```

where *platform* refers to the appropriate hardware platform.

For information about using **installSAM** see: publib.boulder.ibm.com/tividd/td/IBMTivoliSystemAutomationforMultiplatforms3.1.html.

What to do next

Whether you use the DB2 installer or the **installSAM** installation script to install SA MP, follow the same post-installation steps. For more information about general post-install steps, see: Chapter 64, “Installing IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 519

If you use the DB2 High Availability (HA) Feature with IBM Tivoli System Automation for Multiplatforms (SA MP) as your cluster manager, the database manager uses scripts to support automated failover solutions. These scripts are installed or updated automatically when you use the DB2 installer to install or update SA MP. When you install or update SA MP using the **installSAM** utility, you must then manually install or update these scripts. For more information about installing or upgrading the scripts manually, see: Chapter 67, “Installing,

updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 531.

Chapter 65. Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP)

You can upgrade IBM Tivoli System Automation for Multiplatforms (SA MP) using either the DB2 installer or the **installSAM** install script that is included in the IBM DB2 server install media.

Before you begin

Whether you are using the DB2 installer, **installSAM**, or **uninstallSAM**, you must meet the basic prerequisites for installing, upgrading, or uninstalling SA MP. See Part 11, “Installing and upgrading SA MP with the DB2 installer,” on page 515.

If you already have SA MP installed, you can upgrade the installed version of SA MP using the DB2 installer or the **installSAM** installation script. For more information about upgrading SA MP, see Chapter 65, “Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP).”

If a IBM PowerHA SystemMirror for AIX cluster is running, you cannot perform a SA MP installation or upgrade because SA MP bundles Reliable Scalable Cluster Technology (RSCT) filesets that are dependent on PowerHA SystemMirror. To skip the SA MP installation use the **db2_install** command or the **installFixPack** command. For information about installing or upgrading SA MP using a PowerHA SystemMirror cluster, see the white paper entitled Upgrade guide for “DB2 Servers in HACMP Environments”, which is available from the IBM Support and downloads website (<http://www.ibm.com/support/docview.wss?uid=swg21045033>).

About this task

Restrictions

- You cannot upgrade SA MP using the DB2 installer if you have one or more IBM Reliable Scalable Cluster Technology (RSCT) peer domains defined on your system.
- Node-by-node migration is not supported for Version 2.2 to Version 3.2 upgrades. You need to migrate the entire domain. For more information, see the IBM Tivoli System Automation for Multiplatforms Installation and Configuration Guide Version 3.1 (SC34-2584-01) and go to Chapter 1 “Installing and upgrading IBM Tivoli System Automation”, section “Migrating IBM Tivoli System Automation”.

Procedure

There are two methods for installing or upgrading SA MP:

- Using the DB2 installer.
- Using the **installSAM** install script that is included in the IBM Data Server install media.

What to do next

Look in the SA MP installation log for diagnostic information about any warnings or errors that the DB2 installer or the **installSAM** installation script returned. For

more information about the SA MP installation log, see Chapter 68, “IBM Tivoli System Automation for Multiplatforms (SA MP) install and uninstall logs,” on page 535.

Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP) using the DB2 installer

You can upgrade IBM Tivoli System Automation for Multiplatforms (SA MP) using the DB2 installer.

Before you begin

Whether you use the DB2 installer or the **installSAM** installation script that is included in the DB2 installation media to upgrade SA MP, you must meet the basic prerequisites for upgrading SA MP. See Chapter 65, “Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 523.

About this task

There are three methods for using the DB2 installer to install or uninstall:

- DB2 Setup wizard (install, upgrade, or uninstall)
- Silent installation by using a response file with **db2setup** (install or upgrade) or **db2unins** (for uninstall)
- **db2_install** command (for install), **installFixPack** command (for upgrade), or **db2_deinstall** command (for uninstall)

Before upgrading SA MP on a given machine, the DB2 installer queries your system for the following information:

- If SA MP is already installed, is the version of SA MP that is already installed older than the version of SA MP that is on the DB2 installation media?

The DB2 installer calls the **installSAM** installation script to perform some parts of the SA MP upgrade operation. You can call **installSAM** directly. For more information about using the **installSAM** script to upgrade SA MP, see “Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP) using the installSAM installation script” on page 525.

You can use the **-l** option with **db2setup**, **db2_install**, or **installFixPack** to specify where the **installSAM** script saves the SA MP installation log. For more information about the SA MP installation log, see: Chapter 68, “IBM Tivoli System Automation for Multiplatforms (SA MP) install and uninstall logs,” on page 535.

Procedure

- To upgrade SA MP using DB2 Setup wizard, run DB2 Setup wizard and follow the instructions in the wizard.

The information about your system that the DB2 installer collects determines which panels are displayed by the graphical interface of the DB2 Setup wizard during the upgrade. For example, if the version of SA MP that you already have installed is at the same version as, or at a later version than the version of SA MP that is on the DB2 installation media, then the DB2 Setup wizard will not display a panel to upgrade SA MP.

- To upgrade SA MP using a response file, set the response file keyword **INSTALL_TYPE** to "CUSTOM" and then set **COMP** to "TSAMP".

In a response file installation, the default behavior of the DB2 installer is to upgrade SA MP if the version of SA MP that is already installed is older than the version that is on the DB2 installation media. If `INSTALL_TSAMP` is "YES", or if `INSTALL_TSAMP` is commented out or missing from the response file, the DB2 installer upgrades SA MP.

To prevent the DB2 installer from upgrading SA MP in a response file installation, set `INSTALL_TSAMP` to "NO".

Important: The `INSTALL_TSAMP` response file keyword has been deprecated in Version 10.1 and might be removed in a future release. Starting in Version 10.1, IBM Tivoli System Automation for Multiplatforms (SA MP) is automatically installed when SA MP is a required component for other components selected for installation. For more information, see "IBM Tivoli System Automation for Multiplatforms (SA MP) is now automatically installed" in *What's New for DB2 Version 10.1*.

- To upgrade SA MP using the `db2_install` command or the `installFixPack` command, you can run the commands without any parameters specific to SA MP.

The default behavior is to upgrade SA MP if the version of SA MP that is already installed is older than the version that is on the DB2 install media.

To prevent the SA MP upgrade, use the `-f NOTSAMP` option.

What to do next

Whether you use the DB2 installer or the `installSAM` script to upgrade SA MP, follow the same post-upgrade steps. For more information about general post-upgrade steps, see: Chapter 64, "Installing IBM Tivoli System Automation for Multiplatforms (SA MP)," on page 519

Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP) using the `installSAM` installation script

You can upgrade IBM Tivoli System Automation for Multiplatforms (SA MP) using the `installSAM` installation script that is included in the DB2 installation media.

Before you begin

Whether you use the DB2 installer or the `installSAM` installation script that is included in the DB2 installation media to upgrade SA MP, you must meet the basic prerequisites for upgrading SA MP. See Chapter 65, "Upgrading IBM Tivoli System Automation for Multiplatforms (SA MP)," on page 523.

Procedure

Run the `installSAM` installation script

The `installSAM` script is located on the DB2 installation media at the following location:

`db2/platform/tsamp`

Where *platform* refers to the appropriate hardware platform.

For information about running `installSAM`, see: publib.boulder.ibm.com/tividd/td/IBMTivoliSystemAutomationforMultiplatforms3.1.html.

What to do next

Whether you use the DB2 installer or the **installSAM** installation script to install SA MP, follow the same post-installation steps. For more information about general post-install steps, see: Chapter 64, “Installing IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 519

If you use the DB2 High Availability (HA) Feature with IBM Tivoli System Automation for Multiplatforms (SA MP) as your cluster manager, the database manager uses scripts to support automated failover solutions. These scripts are installed or updated automatically when you use the DB2 installer to install or update SA MP. When you install or update SA MP using the **installSAM** utility, you must then manually install or update these scripts. For more information about installing or upgrading the scripts manually, see: Chapter 67, “Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 531.

Chapter 66. Uninstalling IBM Tivoli System Automation for Multiplatforms (SA MP)

You can uninstall IBM Tivoli System Automation for Multiplatforms (SA MP) using either the DB2 installer or the **uninstallSAM** uninstall script that is included in the IBM Data Server install media.

Before you begin

Whether you are using the DB2 installer, **installSAM**, or **uninstallSAM**, you must meet the basic prerequisites for installing, upgrading, or uninstalling SA MP. See Part 11, “Installing and upgrading SA MP with the DB2 installer,” on page 515.

Procedure

There are two methods for uninstalling SA MP:

- Using the DB2 installer
- Using the **uninstallSAM** uninstall script that is included in the IBM Data Server install media

What to do next

Look in the SA MP uninstall log for diagnostic information about any warnings or errors that the DB2 installer or the **uninstallSAM** uninstall script returned. For more information about the SA MP uninstall log, see: Chapter 68, “IBM Tivoli System Automation for Multiplatforms (SA MP) install and uninstall logs,” on page 535.

Uninstalling IBM Tivoli System Automation for Multiplatforms (SA MP) using the DB2 installer

You can uninstall IBM Tivoli System Automation for Multiplatforms (SA MP) using the DB2 installer.

Before you begin

Whether you use the DB2 installer or the **uninstallSAM** uninstall script that is included in the IBM Data Server install media to uninstall SA MP, you must meet the basic prerequisites for uninstalling SA MP. See: Chapter 66, “Uninstalling IBM Tivoli System Automation for Multiplatforms (SA MP).”

About this task

There are three methods for using the DB2 installer to install or uninstall:

- DB2 Setup wizard (install, upgrade, or uninstall)
- Silent installation by using a response file with **db2setup** (install or upgrade) or **db2unins** (for uninstall)
- **db2_install** command (for install), **installFixPack** command (for upgrade), or **db2_deinstall** command (for uninstall)

The DB2 installer calls the **uninstallSAM** install script to perform parts of the SA MP uninstall. You can call **uninstallSAM** directly. For more information about using

the **uninstallSAM** script to uninstall SA MP, see: “Uninstalling IBM Tivoli System Automation for Multiplatforms (SA MP) using the uninstallSAM uninstall script.”

You can use the **-l** option with **db2setup**, **db2_install**, or **installFixPack** to specify where the **installSAM** script saves the SA MP installation log. For more information about the SA MP installation log, see: Chapter 68, “IBM Tivoli System Automation for Multiplatforms (SA MP) install and uninstall logs,” on page 535.

Procedure

To uninstall SA MP using **db2_deinstall**, you can run **db2_deinstall** with the **-a -F TSAMP** option.

By default, the DB2 installer will not uninstall SA MP when you run **db2_deinstall**.

What to do next

Whether you use the DB2 installer or the **uninstallSAM** uninstall script that is included with the IBM Data Server install media to uninstall SA MP, follow the same post-uninstall steps. For more information about general post-uninstall steps, see: Chapter 66, “Uninstalling IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 527

Uninstalling IBM Tivoli System Automation for Multiplatforms (SA MP) using the uninstallSAM uninstall script

You can uninstall IBM Tivoli System Automation for Multiplatforms (SA MP) using the **uninstallSAM** uninstall script that is included in the IBM Data Server install media.

Before you begin

Whether you use the DB2 installer or the **uninstallSAM** uninstall script that is included in the IBM Data Server install media to uninstall SA MP, you must meet the basic prerequisites for uninstalling SA MP. See: Chapter 66, “Uninstalling IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 527.

Procedure

Run the **uninstallSAM** uninstall script

The **uninstallSAM** uninstall script is located on the IBM Data Server media at the following location:

`db2/platform/tsamp`

where *platform* refers to the appropriate hardware platform.

For information about using **uninstallSAM** see: publib.boulder.ibm.com/tividd/td/IBMTivoliSystemAutomationforMultiplatforms3.1.html.

What to do next

Whether you use the DB2 installer or the **installSAM** installation script to install SA MP, follow the same post-installation steps. For more information about general post-install steps, see: Chapter 64, “Installing IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 519

If you use the DB2 High Availability (HA) Feature with IBM Tivoli System Automation for Multiplatforms (SA MP) as your cluster manager, the database manager uses scripts to support automated failover solutions. These scripts are uninstalled automatically when you run **db2_deinstall** to uninstall SA MP. When you uninstall SA MP using the **uninstallSam** utility, then you must manually uninstall these scripts. For more information about uninstalling the scripts manually, see: Chapter 67, “Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 531.

Chapter 67. Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP)

If you use the DB2 High Availability (HA) Feature with IBM Tivoli System Automation for Multiplatforms (SA MP) as your cluster manager, the database manager uses scripts to support automated failover solutions. You can install, upgrade, and uninstall these SA MP scripts using the DB2 installer, or the **installSAM** or **uninstallSAM** scripts that are included in the DB2 install media.

Before you begin

- To install, upgrade, or uninstall the SA MP scripts using the DB2 installer or the **installSAM** or **uninstallSAM** scripts that are included in the DB2 install media, you must purchase the DB2 High Availability Feature.
- You must have root authority to install, upgrade, or uninstall the SA MP scripts. If you perform a non-root install of IBM Data Server, you can install the SA MP scripts from the DB2 install media separately. When you install the SA MP scripts separately, you still must have root authority.

Procedure

There are two methods for installing, upgrading, and uninstalling the SA MP automated failover scripts:

- Using the DB2 installer. Refer to “Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP) using DB2 installer.”
- Manually installing from the DB2 install media. Refer to “Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP) manually” on page 532.

Results

When you install the SA MP automated failover scripts, the scripts are installed in the following location:

```
/usr/sbin/rsct/sapolicies/db2
```

If you uninstall the SA MP scripts, you can no longer use HADR functions within a cluster managed by SA MP.

Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP) using DB2 installer

You can install, upgrade, or uninstall IBM Tivoli System Automation for Multiplatforms (SA MP) automated failover scripts using the DB2 installer.

Before you begin

Whether you use the DB2 installer or manually install, upgrade, or uninstall SA MP automated failover scripts, you must meet the basic prerequisites for installing,

upgrading, and uninstalling SA MP automated failover scripts. See: Chapter 67, “Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 531.

About this task

There are three methods for using the DB2 installer to install or uninstall:

- DB2 Setup wizard (install, upgrade, or uninstall)
- Silent installation by using a response file with **db2setup** (install or upgrade) or **db2unins** (for uninstall)
- **db2_install** command (for install), **installFixPack** command (for upgrade), or **db2_deinstall** command (for uninstall)

Procedure

1. To install the SA MP automated failover scripts, run the DB2 installer.
By default, the DB2 installer installs the SA MP automated failover scripts if SA MP is installed or being installed, and the scripts are not already installed.
2. To upgrade the SA MP automated failover scripts, run the DB2 installer.
By default, the DB2 installer will upgrade the SA MP automated failover scripts if SA MP is installed or being installed, and the scripts that are already installed are at a lower version than the version of the scripts that are on the DB2 install media.
3. To uninstall the SA MP automated failover scripts, run the DB2 installer.

Results

Whether you use the DB2 installer or manually install, upgrade, or uninstall the SA MP automated failover scripts, the general results are the same, see Chapter 67, “Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 531.

Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP) manually

You can manually install, upgrade, or uninstall IBM Tivoli System Automation for Multiplatforms (SA MP) automated failover scripts from the DB2 install media.

Before you begin

Whether you use the DB2 installer or manually install, upgrade, or uninstall SA MP automated failover scripts, you must meet the basic prerequisites for installing, upgrading, and uninstalling SA MP automated failover scripts. See: Chapter 67, “Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 531.

About this task

The SA MP automated failover scripts are installed automatically by the DB2 installer when SA MP is installed or being installed. If you install or upgrade SA MP manually, then you must install or upgrade the SA MP scripts manually. The DB2 installer does not uninstall the SA MP scripts, so to remove these scripts, you must uninstall them manually.

Procedure

To install, update, or uninstall the SA MP automated failover scripts manually, use the **db2cptsa** utility.

Results

Whether you use the DB2 installer or manually install, upgrade, or uninstall the SA MP automated failover scripts, the general results are the same, see Chapter 67, “Installing, updating, and uninstalling automated failover scripts for the IBM Tivoli System Automation for Multiplatforms (SA MP),” on page 531.

Chapter 68. IBM Tivoli System Automation for Multiplatforms (SA MP) install and uninstall logs

Diagnostic information, warnings, and error messages related to installing, upgrading, and uninstalling IBM Tivoli System Automation for Multiplatforms (SA MP) are located in install and uninstall logs specific to SA MP.

You can install, upgrade, or uninstall SA MP using the DB2 installer, or the **installSAM** or **uninstallSAM** scripts that are included on the IBM Data Server install media. In fact, the DB2 installer uses the **installSAM** utility and the **uninstallSAM** utility to perform part of the install, upgrade, and uninstall operations.

The **installSAM** utility generates a sequence of log files that are named sequentially:

```
/tmp/installSAM.log-number.log
```

where *log-number* identifies the log file in the sequence.

You can use the **-l** option with **db2setup**, **db2_install**, or **installFixPack** to specify where the **installSAM** utility should place the SA MP install log.

The **uninstallSAM** utility generates a sequence of log files that are named sequentially:

```
/tmp/uninstallSAM.log-number.log
```

log-number identifies the log file in the sequence.

You can use the **-l** option with **db2unins** or **db2_deinstall** to specify where the **uninstallSAM** utility should place the SA MP uninstall log.

Chapter 69. License terms for using IBM Tivoli System Automation for Multiplatforms (SA MP) integrated with IBM DB2 server

There are conditions under which you can use IBM Tivoli System Automation for Multiplatforms (SA MP) that is integrated with IBM DB2 server.

You can use SA MP that is integrated with DB2 server with DB2 high availability disaster recovery (HADR) functionality if you have purchased a license for one of:

- DB2 Advanced Enterprise Server Edition
- DB2 Enterprise Server Edition
- DB2 Connect Enterprise Edition
- DB2 Workgroup Server Edition

Also you can use SA MP that is integrated with DB2 server with HADR functionality if you have purchased a license for the DB2 High Availability Feature and one of:

- DB2 Express Edition

You can use a try and buy version of SA MP that is integrated with DB2 server with HADR functionality if you have a try and buy license for one of:

- DB2 Advanced Enterprise Server Edition
- DB2 Enterprise Server Edition
- DB2 Connect Enterprise Edition
- DB2 Workgroup Server Edition
- DB2 Express Edition

Chapter 70. Supported software and hardware for IBM Tivoli System Automation for Multiplatforms (SA MP)

IBM Tivoli System Automation for Multiplatforms (SA MP) is supported on AIX, Linux, Solaris SPARC, and Windows. SA MP is integrated with IBM DB2 server on AIX, Linux, and Solaris SPARC. On Windows, SA MP is bundled separately as part of DB2 server.

SA MP is integrated or bundled with the following DB2 database products and features:

- DB2 Advanced Enterprise Server Edition
- DB2 Enterprise Server Edition
- DB2 Connect Enterprise Edition
- DB2 Workgroup Server Edition
- DB2 Express-C with Fixed Term License (FTL)
- IBM DB2 High Availability Feature for Express Edition.

The following table shows which versions of SA MP are included in the DB2 installation media or in the fix pack images. For more information about the operating systems supported, see the system requirements information.

Table 54. SA MP versions shipped in DB2 Version 10.1

Operating system platform	SA MP Version installed
AIX	<ul style="list-style-type: none">• DB2 Version 10.1 : SA MP 3.2.2.1
Linux distributions	<ul style="list-style-type: none">• DB2 Version 10.1: SA MP 3.2.2.1 <p>Note:</p> <ul style="list-style-type: none">• SA MP is not supported on the Ubuntu distribution.
HP-UX	<ul style="list-style-type: none">• SA MP is not supported.
Solaris 10 on SPARC (64-bit)	<ul style="list-style-type: none">• SA MP 3.2.2.1
Solaris on x64	<ul style="list-style-type: none">• SA MP is not supported.
Windows	<ul style="list-style-type: none">• SA MP 3.2.2 <p>Note: This version is bundled with the DB2 Version 10.1 GA installation media and not integrated with the DB2 product installation.</p>

Note: SA MP Version 3.1 and Version 3.2 are not supported on AIX system workload partitions (WPARs), as well as on domains with nodes that are Solaris zone containers.

Table 55. Versions of SA MP supported in DB2 Version 9.7

Operating system platform	SA MP product version supported
AIX	<ul style="list-style-type: none">• SA MP 3.1.x• SA MP 3.2.x

Table 55. Versions of SA MP supported in DB2 Version 9.7 (continued)

Operating system platform	SA MP product version supported
Linux distributions	<ul style="list-style-type: none"> • SA MP 3.1.x • SA MP 3.2.x
HP-UX	<ul style="list-style-type: none"> • SA MP is not supported.
Solaris 10 on SPARC (64-bit)	<ul style="list-style-type: none"> • SA MP 3.1.x • SA MP 3.2.x
Solaris on x64	<ul style="list-style-type: none"> • SA MP is not supported.
Windows	<ul style="list-style-type: none"> • SA MP 3.1.x • SA MP 3.2.x

Note: where *x* is any SA MP modification level and fix pack.

If you plan on using the DB2 High Availability Feature, ensure that your system meets the prerequisites for IBM Tivoli System Automation for Multiplatforms (SA MP). For more details, see the Installation and Configuration guides at Tivoli software information center.

If you are not using the integrated or bundled copy of IBM Tivoli System Automation for Multiplatforms (SA MP), refer to the following website for the full list of supported operating systems: <http://www.ibm.com/software/tivoli/products/sys-auto-linux/platforms.html>.

Part 12. IBM Data Studio integrated installation with DB2 database products

IBM Data Studio installation has been integrated with the DB2 product installation. The DB2 Setup wizard now provides an option to install IBM Data Studio components.

Chapter 71. Installing IBM Data Studio using the DB2 Setup wizard (Windows)

Start the DB2 Setup wizard to define your installation preferences and to install IBM Data Studio components.

Before you begin

Before you start the DB2 Setup wizard:

- Ensure that the computer where you are installing IBM Data Studio meets the installation, memory, disk, and user privilege requirements. For details, see Installation requirements in the IBM Data Studio Information Center.
- Ensure that the computer where you are installing IBM Data Studio has an internet connection. If an internet connection is not possible, you must download this tool to the computer.
- You must have a local Administrator user account with the recommended user rights to perform the installation.

Note: If a non-Administrator user account is going to do the product installation, then the VS2010 runtime library must be installed before attempting to install a DB2 database product. The VS2010 runtime library is needed on the operating system before the DB2 database product can be installed. The VS2010 runtime library is available from the Microsoft runtime library download website. There are two choices: choose `vcredist_x86.exe` for 32-bit systems or `vcredist_x64.exe` for 64-bit systems.

- You should close all programs so that the installation program can update any files on the computer without requiring a reboot.
- For installations from a virtual drive, you must map the network drive to a Windows drive letter. The DB2 Setup wizard does not support installation from a virtual drive or an unmapped network drive (such as `\\hostname\sharename` in Windows Explorer).

About this task

If you downloaded IBM Data Studio, perform the Installing IBM Data Studio full client with the Installation Manager Install wizard task in the IBM Data Studio Information Center.

Restrictions

- The DB2 Setup wizard fields do not accept non-English characters.
- If you enable extended security on Windows Vista or Windows 2008, or higher, users must belong to the DB2ADMNS or DB2USERS group to run local DB2 commands and applications because of an extra security feature (User Access Control) that limits the privileges that local administrators have by default. If users do not belong to one of these groups, they will not have read access to local DB2 configuration or application data.

Procedure

To install the IBM Data Studio using the DB2 Setup wizard:

1. Log on to the system with the local Administrator account that you have defined for the DB2 installation.
2. If you have the DB2 database product DVD, insert it into the drive. If enabled, the autorun feature automatically starts the DB2 Setup Launchpad. If the autorun does not work, use Windows Explorer to browse the DB2 database product DVD and double-click the **setup** icon to start the DB2 Setup Launchpad.
3. If you downloaded the DB2 database product from Passport Advantage, run the executable file to extract the DB2 database product installation files. Use Windows Explorer to browse the DB2 installation files and double-click the **setup** icon to start the DB2 Setup Launchpad.
4. From the DB2 Setup launchpad, you can view installation prerequisites and the release notes, or you can proceed directly to the installation. You might want to review the installation prerequisites and release notes for late-breaking information.
5. Click **Install Data Studio Components** and the Install Data Studio Components window displays the products available for installation.
6. To install the IBM Data Studio client, click **Install**. The IBM Installation Manager is called and downloads the IBM Data Studio client if necessary. The Install Packages window displays the packages that you can install. Select the packages that you want to install and click **Next** to proceed with the installation.
7. To download and install the IBM Data Studio web console, click **Download IBM Data Studio**. This opens a browser window and directs you to the download web page. Follow the instructions to download and install IBM Data Studio web console.

Chapter 72. Installing IBM Data Studio using the DB2 Setup wizard (Linux)

Start the DB2 Setup wizard to define your installation preferences and to install IBM Data Studio components.

Before you begin

Before you start the DB2 Setup wizard:

- Ensure that the computer where you are installing IBM Data Studio meets the installation, memory, disk, and user privilege requirements. For details, see Installation requirements in the IBM Data Studio Information Center.
- Ensure that the computer where you are installing IBM Data Studio has an internet connection. If an internet connection is not possible, you must download this tool to the computer.
- Ensure you have a supported browser installed.
- Ensure that the DB2 database product image is available on the computer. You can obtain a DB2 installation image either by purchasing a physical DB2 database product DVD, or by downloading an installation image from Passport Advantage.
- If you are installing a non-English version of a DB2 database product, you must have the appropriate National Language Packages.
- Ensure that you have installed X Linux software capable of rendering a graphical user, the X Linux server is running, the *DISPLAY* variable is defined. The DB2 Setup wizard is a graphical installer.
- If you are using security software in your environment, you must manually create required DB2 users before you start the DB2 Setup wizard.

Restrictions

- You cannot have more than one instance of the DB2 Setup wizard running in any user account.
- The DB2 Setup wizard fields do not accept non-English characters.

Procedure

To install the IBM Data Studio using the DB2 Setup wizard:

1. If you have a physical DB2 database product DVD, change to the directory where the DB2 database product DVD is mounted by entering the following command:

```
cd /dvdrom
```

where */dvdrom* represents the mount point of the DB2 database product DVD.

2. If you downloaded the DB2 database product image, you must extract and untar the product file.

- a. Extract the product file:

```
gzip -d product.tar.gz
```

where *product* is the name of the product that you downloaded.

- b. Untar the product file:

```
tar -xvf product.tar
```

where *product* is the name of the product that you downloaded.

c. Change directory:

```
cd ./product
```

where *product* is the name of the product that you downloaded.

Note: If you downloaded a National Language Package, untar it into the same directory. This will create the subdirectories (for example *./nlpack*) in the same directory, and allows the installer to automatically find the installation images without prompting.

3. Enter the ***./db2setup*** command from the directory where the database product image resides to start the DB2 Setup wizard.
4. Click **Install Data Studio Components** and the Install Data Studio Components window displays the products available for installation.
5. To download and install the IBM Data Studio client or the IBM Data Studio web console, click **Download IBM Data Studio**. This opens a browser window and directs you to the download web page. Follow the instructions to download and install IBM Data Studio client or IBM Data Studio web console.

Chapter 73. Scenario: Installation of IBM Data Studio client

This scenario describes the installation of IBM Data Studio client from the DB2 Enterprise Server Edition product media.

The computer is running Linux operating system and has an internet connection.

Starting the DB2 Setup wizard

To start the DB2 Setup wizard, run the **db2setup** command from the directory where the DB2 installation image resides as follows:

```
./db2setup -l /tmp/db2setup.log -t /tmp/db2setup.trc
```

The **-l** and **-t** recommended parameters enable DB2 logging and tracing.

After you issue this command, the IBM DB2 Setup Launchpad opens. The following figure shows the DB2 Setup Launchpad:

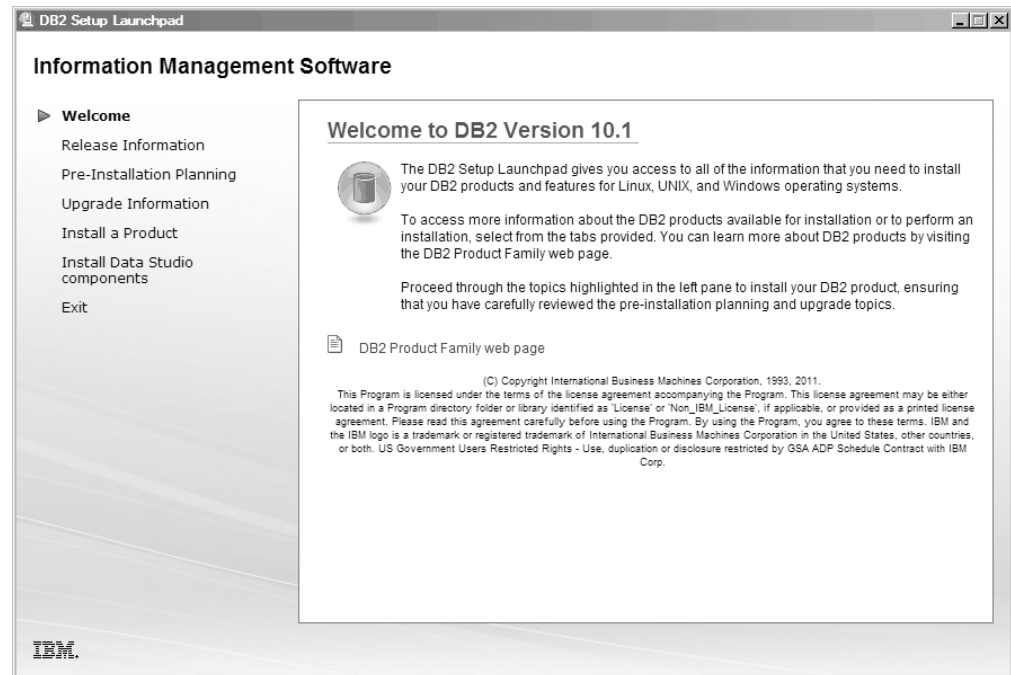


Figure 46. DB2 Setup launchpad

From this window, you can view the DB2 Enterprise Server Edition release information, upgrade options and pre-installation planning, or you can proceed directly to the IBM Data Studio installation.

Installing IBM Data Studio

To install the IBM Data Studio components, click **Install Data Studio components**. The Install Data Studio components panel displays the products available for installation.

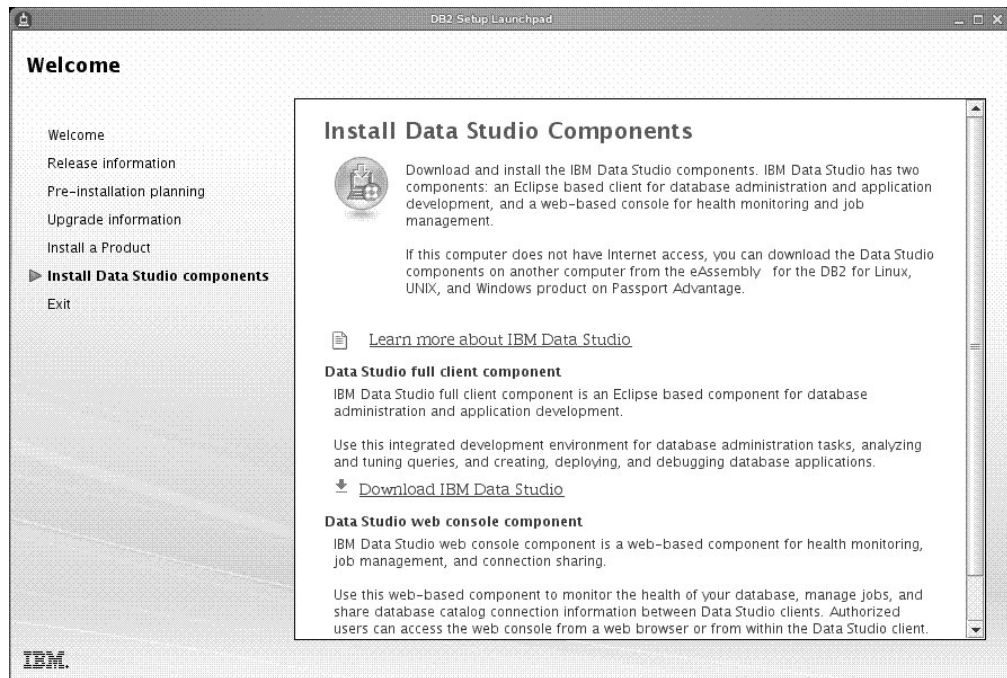


Figure 47. Install Data Studio components Panel

Click **Download IBM Data Studio** to download the IBM Data Studio client. This opens a browser window and directs you to the download web page.

Follow the installation instructions specified in the release notes.

Part 13. Appendixes

Appendix A. Overview of the DB2 technical information

DB2 technical information is available in multiple formats that can be accessed in multiple ways.

DB2 technical information is available through the following tools and methods:

- DB2 Information Center
 - Topics (Task, concept and reference topics)
 - Sample programs
 - Tutorials
- DB2 books
 - PDF files (downloadable)
 - PDF files (from the DB2 PDF DVD)
 - printed books
- Command-line help
 - Command help
 - Message help

Note: The DB2 Information Center topics are updated more frequently than either the PDF or the hardcopy books. To get the most current information, install the documentation updates as they become available, or refer to the DB2 Information Center at ibm.com.

You can access additional DB2 technical information such as technotes, white papers, and IBM Redbooks® publications online at ibm.com. Access the DB2 Information Management software library site at <http://www.ibm.com/software/data/sw-library/>.

Documentation feedback

We value your feedback on the DB2 documentation. If you have suggestions for how to improve the DB2 documentation, send an email to db2docs@ca.ibm.com. The DB2 documentation team reads all of your feedback, but cannot respond to you directly. Provide specific examples wherever possible so that we can better understand your concerns. If you are providing feedback on a specific topic or help file, include the topic title and URL.

Do not use this email address to contact DB2 Customer Support. If you have a DB2 technical issue that the documentation does not resolve, contact your local IBM service center for assistance.

DB2 technical library in hardcopy or PDF format

The following tables describe the DB2 library available from the IBM Publications Center at www.ibm.com/e-business/linkweb/publications/servlet/pbi.wss. English and translated DB2 Version 10.1 manuals in PDF format can be downloaded from www.ibm.com/support/docview.wss?rs=71&uid=swg27009474.

Although the tables identify books available in print, the books might not be available in your country or region.

The form number increases each time a manual is updated. Ensure that you are reading the most recent version of the manuals, as listed below.

Note: The *DB2 Information Center* is updated more frequently than either the PDF or the hard-copy books.

Table 56. DB2 technical information

Name	Form Number	Available in print	Last updated
<i>Administrative API Reference</i>	SC27-3864-00	Yes	April, 2012
<i>Administrative Routines and Views</i>	SC27-3865-01	No	January, 2013
<i>Call Level Interface Guide and Reference Volume 1</i>	SC27-3866-01	Yes	January, 2013
<i>Call Level Interface Guide and Reference Volume 2</i>	SC27-3867-01	Yes	January, 2013
<i>Command Reference</i>	SC27-3868-01	Yes	January, 2013
<i>Database Administration Concepts and Configuration Reference</i>	SC27-3871-01	Yes	January, 2013
<i>Data Movement Utilities Guide and Reference</i>	SC27-3869-01	Yes	January, 2013
<i>Database Monitoring Guide and Reference</i>	SC27-3887-01	Yes	January, 2013
<i>Data Recovery and High Availability Guide and Reference</i>	SC27-3870-01	Yes	January, 2013
<i>Database Security Guide</i>	SC27-3872-01	Yes	January, 2013
<i>DB2 Workload Management Guide and Reference</i>	SC27-3891-01	Yes	January, 2013
<i>Developing ADO.NET and OLE DB Applications</i>	SC27-3873-01	Yes	January, 2013
<i>Developing Embedded SQL Applications</i>	SC27-3874-01	Yes	January, 2013
<i>Developing Java Applications</i>	SC27-3875-01	Yes	January, 2013
<i>Developing Perl, PHP, Python, and Ruby on Rails Applications</i>	SC27-3876-00	No	April, 2012
<i>Developing RDF Applications for IBM Data Servers</i>	SC27-4462-00	Yes	January, 2013
<i>Developing User-defined Routines (SQL and External)</i>	SC27-3877-01	Yes	January, 2013
<i>Getting Started with Database Application Development</i>	GI13-2046-01	Yes	January, 2013

Table 56. DB2 technical information (continued)

Name	Form Number	Available in print	Last updated
<i>Getting Started with DB2 Installation and Administration on Linux and Windows</i>	GI13-2047-00	Yes	April, 2012
<i>Globalization Guide</i>	SC27-3878-00	Yes	April, 2012
<i>Installing DB2 Servers</i>	GC27-3884-01	Yes	January, 2013
<i>Installing IBM Data Server Clients</i>	GC27-3883-00	No	April, 2012
<i>Message Reference Volume 1</i>	SC27-3879-01	No	January, 2013
<i>Message Reference Volume 2</i>	SC27-3880-01	No	January, 2013
<i>Net Search Extender Administration and User's Guide</i>	SC27-3895-01	No	January, 2013
<i>Partitioning and Clustering Guide</i>	SC27-3882-01	Yes	January, 2013
<i>Preparation Guide for DB2 10.1 Fundamentals Exam 610</i>	SC27-4540-00	No	January, 2013
<i>Preparation Guide for DB2 10.1 DBA for Linux, UNIX, and Windows Exam 611</i>	SC27-4541-00	No	January, 2013
<i>pureXML Guide</i>	SC27-3892-01	Yes	January, 2013
<i>Spatial Extender User's Guide and Reference</i>	SC27-3894-00	No	April, 2012
<i>SQL Procedural Languages: Application Enablement and Support</i>	SC27-3896-01	Yes	January, 2013
<i>SQL Reference Volume 1</i>	SC27-3885-01	Yes	January, 2013
<i>SQL Reference Volume 2</i>	SC27-3886-01	Yes	January, 2013
<i>Text Search Guide</i>	SC27-3888-01	Yes	January, 2013
<i>Troubleshooting and Tuning Database Performance</i>	SC27-3889-01	Yes	January, 2013
<i>Upgrading to DB2 Version 10.1</i>	SC27-3881-01	Yes	January, 2013
<i>What's New for DB2 Version 10.1</i>	SC27-3890-01	Yes	January, 2013
<i>XQuery Reference</i>	SC27-3893-01	No	January, 2013

Table 57. DB2 Connect-specific technical information

Name	Form Number	Available in print	Last updated
<i>DB2 Connect Installing and Configuring DB2 Connect Personal Edition</i>	SC27-3861-00	Yes	April, 2012

Table 57. DB2 Connect-specific technical information (continued)

Name	Form Number	Available in print	Last updated
DB2 Connect Installing and Configuring DB2 Connect Servers	SC27-3862-01	Yes	January, 2013
DB2 Connect User's Guide	SC27-3863-01	Yes	January, 2013

Displaying SQL state help from the command line processor

DB2 products return an SQLSTATE value for conditions that can be the result of an SQL statement. SQLSTATE help explains the meanings of SQL states and SQL state class codes.

Procedure

To start SQL state help, open the command line processor and enter:

`? sqlstate` or `? class code`

where *sqlstate* represents a valid five-digit SQL state and *class code* represents the first two digits of the SQL state.

For example, `? 08003` displays help for the 08003 SQL state, and `? 08` displays help for the 08 class code.

Accessing different versions of the DB2 Information Center

Documentation for other versions of DB2 products is found in separate information centers on ibm.com.

About this task

For DB2 Version 10.1 topics, the *DB2 Information Center* URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v10r1>.

For DB2 Version 9.8 topics, the *DB2 Information Center* URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9r8/>.

For DB2 Version 9.7 topics, the *DB2 Information Center* URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9r7/>.

For DB2 Version 9.5 topics, the *DB2 Information Center* URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9r5>.

For DB2 Version 9.1 topics, the *DB2 Information Center* URL is <http://publib.boulder.ibm.com/infocenter/db2luw/v9/>.

For DB2 Version 8 topics, go to the *DB2 Information Center* URL at: <http://publib.boulder.ibm.com/infocenter/db2luw/v8/>.

Updating the DB2 Information Center installed on your computer or intranet server

A locally installed DB2 Information Center must be updated periodically.

Before you begin

A DB2 Version 10.1 Information Center must already be installed. For details, see the “Installing the DB2 Information Center using the DB2 Setup wizard” topic in *Installing DB2 Servers*. All prerequisites and restrictions that applied to installing the Information Center also apply to updating the Information Center.

About this task

An existing DB2 Information Center can be updated automatically or manually:

- Automatic updates update existing Information Center features and languages. One benefit of automatic updates is that the Information Center is unavailable for a shorter time compared to during a manual update. In addition, automatic updates can be set to run as part of other batch jobs that run periodically.
- Manual updates can be used to update existing Information Center features and languages. Automatic updates reduce the downtime during the update process, however you must use the manual process when you want to add features or languages. For example, a local Information Center was originally installed with both English and French languages, and now you want to also install the German language; a manual update will install German, as well as, update the existing Information Center features and languages. However, a manual update requires you to manually stop, update, and restart the Information Center. The Information Center is unavailable during the entire update process. In the automatic update process the Information Center incurs an outage to restart the Information Center after the update only.

This topic details the process for automatic updates. For manual update instructions, see the “Manually updating the DB2 Information Center installed on your computer or intranet server” topic.

Procedure

To automatically update the DB2 Information Center installed on your computer or intranet server:

1. On Linux operating systems,
 - a. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the `/opt/ibm/db2ic/V10.1` directory.
 - b. Navigate from the installation directory to the `doc/bin` directory.
 - c. Run the `update-ic` script:
`update-ic`
2. On Windows operating systems,
 - a. Open a command window.
 - b. Navigate to the path where the Information Center is installed. By default, the DB2 Information Center is installed in the `<Program Files>\IBM\DB2 Information Center\Version 10.1` directory, where `<Program Files>` represents the location of the Program Files directory.
 - c. Navigate from the installation directory to the `doc\bin` directory.
 - d. Run the `update-ic.bat` file:
`update-ic.bat`

Results

The DB2 Information Center restarts automatically. If updates were available, the Information Center displays the new and updated topics. If Information Center updates were not available, a message is added to the log. The log file is located in `doc\eclipse\configuration` directory. The log file name is a randomly generated number. For example, `1239053440785.log`.

Manually updating the DB2 Information Center installed on your computer or intranet server

If you have installed the DB2 Information Center locally, you can obtain and install documentation updates from IBM.

About this task

Updating your locally installed *DB2 Information Center* manually requires that you:

1. Stop the *DB2 Information Center* on your computer, and restart the Information Center in stand-alone mode. Running the Information Center in stand-alone mode prevents other users on your network from accessing the Information Center, and allows you to apply updates. The Workstation version of the DB2 Information Center always runs in stand-alone mode. .
2. Use the Update feature to see what updates are available. If there are updates that you must install, you can use the Update feature to obtain and install them

Note: If your environment requires installing the *DB2 Information Center* updates on a machine that is not connected to the internet, mirror the update site to a local file system by using a machine that is connected to the internet and has the *DB2 Information Center* installed. If many users on your network will be installing the documentation updates, you can reduce the time required for individuals to perform the updates by also mirroring the update site locally and creating a proxy for the update site.

If update packages are available, use the Update feature to get the packages. However, the Update feature is only available in stand-alone mode.

3. Stop the stand-alone Information Center, and restart the *DB2 Information Center* on your computer.

Note: On Windows 2008, Windows Vista (and higher), the commands listed later in this section must be run as an administrator. To open a command prompt or graphical tool with full administrator privileges, right-click the shortcut and then select **Run as administrator**.


Procedure

To update the *DB2 Information Center* installed on your computer or intranet server:

1. Stop the *DB2 Information Center*.
 - On Windows, click **Start > Control Panel > Administrative Tools > Services**. Then right-click **DB2 Information Center** service and select **Stop**.
 - On Linux, enter the following command:
`/etc/init.d/db2icdv10 stop`
2. Start the Information Center in stand-alone mode.
 - On Windows:
 - a. Open a command window.

- b. Navigate to the path where the Information Center is installed. By default, the *DB2 Information Center* is installed in the *Program_Files\IBM\DB2 Information Center\Version 10.1* directory, where *Program_Files* represents the location of the Program Files directory.
- c. Navigate from the installation directory to the *doc\bin* directory.
- d. Run the *help_start.bat* file:
`help_start.bat`
- On Linux:
 - a. Navigate to the path where the Information Center is installed. By default, the *DB2 Information Center* is installed in the */opt/ibm/db2ic/V10.1* directory.
 - b. Navigate from the installation directory to the *doc/bin* directory.
 - c. Run the *help_start* script:
`help_start`

The system's default Web browser opens to display the stand-alone Information Center.

3. Click the **Update** button (). (JavaScript must be enabled in your browser.) On the right panel of the Information Center, click **Find Updates**. A list of updates for existing documentation displays.
4. To initiate the installation process, check that the selections you want to install, then click **Install Updates**.
5. After the installation process has completed, click **Finish**.
6. Stop the stand-alone Information Center:
 - On Windows, navigate to the *doc\bin* directory within the installation directory, and run the *help_end.bat* file:
`help_end.bat`
 - Note:** The *help_end* batch file contains the commands required to safely stop the processes that were started with the *help_start* batch file. Do not use Ctrl-C or any other method to stop *help_start.bat*.
 - On Linux, navigate to the *doc/bin* directory within the installation directory, and run the *help_end* script:
`help_end`
 - Note:** The *help_end* script contains the commands required to safely stop the processes that were started with the *help_start* script. Do not use any other method to stop the *help_start* script.
7. Restart the *DB2 Information Center*.
 - On Windows, click **Start > Control Panel > Administrative Tools > Services**. Then right-click **DB2 Information Center** service and select **Start**.
 - On Linux, enter the following command:
`/etc/init.d/db2icdv10 start`

Results

The updated *DB2 Information Center* displays the new and updated topics.

DB2 tutorials

The DB2 tutorials help you learn about various aspects of DB2 database products. Lessons provide step-by-step instructions.

Before you begin

You can view the XHTML version of the tutorial from the Information Center at <http://publib.boulder.ibm.com/infocenter/db2luw/v10r1/>.

Some lessons use sample data or code. See the tutorial for a description of any prerequisites for its specific tasks.

DB2 tutorials

To view the tutorial, click the title.

"pureXML®" in pureXML Guide

Set up a DB2 database to store XML data and to perform basic operations with the native XML data store.

DB2 troubleshooting information

A wide variety of troubleshooting and problem determination information is available to assist you in using DB2 database products.

DB2 documentation

Troubleshooting information can be found in the *Troubleshooting and Tuning Database Performance* or the Database fundamentals section of the *DB2 Information Center*, which contains:

- Information about how to isolate and identify problems with DB2 diagnostic tools and utilities.
- Solutions to some of the most common problem.
- Advice to help solve other problems you might encounter with your DB2 database products.

IBM Support Portal

See the IBM Support Portal if you are experiencing problems and want help finding possible causes and solutions. The Technical Support site has links to the latest DB2 publications, TechNotes, Authorized Program Analysis Reports (APARs or bug fixes), fix packs, and other resources. You can search through this knowledge base to find possible solutions to your problems.

Access the IBM Support Portal at http://www.ibm.com/support/entry/portal/Overview/Software/Information_Management/DB2_for_Linux,_UNIX_and_Windows

Terms and conditions

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability: These terms and conditions are in addition to any terms of use for the IBM website.

Personal use: You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use: You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights: Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Trademarks: IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml

Appendix B. Notices

This information was developed for products and services offered in the U.S.A. Information about non-IBM products is based on information available at the time of first publication of this document and is subject to change.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information about the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country/region where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements, changes, or both in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to websites not owned by IBM are provided for convenience only and do not in any manner serve as an endorsement of those

websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information that has been exchanged, should contact:

IBM Canada Limited
U59/3600
3600 Steeles Avenue East
Markham, Ontario L3R 9Z7
CANADA

Such information may be available, subject to appropriate terms and conditions, including, in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems, and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements, or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility, or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious, and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating

platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. *_enter the year or years_*. All rights reserved.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

The following terms are trademarks or registered trademarks of other companies

- Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
- Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle, its affiliates, or both.
- UNIX is a registered trademark of The Open Group in the United States and other countries.
- Intel, Intel logo, Intel Inside, Intel Inside logo, Celeron, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
- Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Index

Numerics

10GE 208, 221
 setting up on AIX 178

A

AIX
 CD mounting 91
 DB2 home file system creation 328
 distributing commands to multiple nodes 347
 DVD mounting 91
 environment settings 325
 installing
 DB2 products in workload partition 39
 DB2 server products 13, 56
 IBM data server clients 56
 NFS 335
 required users
 creating 336
automatic client reroute
 overview 159

B

BACKUP DATABASE command
 pre-conversion tasks for DB2 servers 447, 452
backups
 databases
 pre-conversion tasks for DB2 servers 447, 452
batch files
 response file installations 385
binding
 database utilities 484

C

CDs
 mounting
 AIX 91
 HP-UX 92
 Linux 93
 Solaris 93
check installation prerequisites tool command 3
client affinities
 overview 159
clients
 DB2 pureScale Feature 159
cluster caching facilities
 adding 280
 adding communication adapter ports 278
cluster file systems
 managed by DB2 product
 cleaning 314
code pages
 supported 82
command line processor (CLP)
 installation verification 119

commands
 db2_deinstall
 removing DB2 products 494
 db2fs
 overview 122
 db2idrop
 removing instances 493
 db2ilist 492
 db2ls
 listing DB2 products and features 139
 db2osconf
 determining kernel configuration parameter values 43
 db2prereqcheck 3
 db2rfe
 enabling root-based features 19
 overcoming limitations on non-root installations 21
 db2sampl
 verifying installation 119
 db2secv82 123
 db2setup
 displaying DB2 Setup wizard in your national
 language 82
 db2stop
 stopping DB2 492
 doce_deinstall
 removing products 494
 remote 357
communications
 fast communication manager (FCM) 34
configuration
 AIX 183
 Linux 206, 216
 PVIDs 188
 RoCE cluster interconnect 184
 setting up IP interfaces on switch on 10GE network 216
 setting up netmon.cf file on 10GE network 229
 switch failover
 AIX 174, 183
 Linux 206, 216, 219
 switches
 RoCE networks with AIX servers 174
contacts
 setting lists 124
conversion
 DB2 pureScale environments
 instances 439, 449
 planning 441
 post-conversion tasks 451
 pre-conversion tasks 443
 planning
 DB2 pureScale Feature 441

D

daemons
 Information Center 429
data server clients
 See IBM data server clients 11
database partition servers
 enabling communications (UNIX) 355

- database partition servers (*continued*)
 - installing using response file
 - Linux 360, 378
 - overview 359
 - UNIX 360, 378
 - Windows 359
- DB2 administration server (DAS)
 - removing 492
 - stopping 491
- DB2 cluster services
 - tiebreaker 151
- DB2 Connect
 - DB2 pureScale Feature 157
- DB2 copies
 - converting instances 435
 - default copy after installation 125
 - managing 137
 - multiple on same computer
 - Linux 90
 - UNIX 90
 - Windows 85
 - uninstalling 501
- DB2 High Availability Feature
 - Tivoli System Automation for Multiplatforms (SA MP) 515
- DB2 Information Center
 - installing
 - DB2 Setup Wizard (Linux) 421
 - DB2 Setup Wizard (Windows) 425
 - fix pack 419
 - Linux 421
 - options 427
 - Windows 425
 - updating 431, 555, 556
 - versions 554
- DB2 installer
 - taking over a user-managed GPFS cluster 274
- DB2 instances
 - converting to new database product 435
- DB2 products
 - listing 139
- DB2 pureScale environment 285
- DB2 pureScale environments
 - db2sshd user ID
 - changing 286
 - re-creating 286
 - monitoring
 - shared file systems 291
 - network topology 145
 - updating instance to DB2 pureScale instance 281
- DB2 pureScale Feature
 - conversion
 - instances 439, 449
 - planning 441
 - post-conversion tasks 451
 - pre-conversion tasks 443
 - db2_install command 267
 - enabling SCSI-3 PR 274
 - GDPCs
 - configuring 293
 - creating databases 309
 - HADR 309
 - installing 294
 - prerequisites 242
 - installing
 - AIX user requirements 170
 - DB2 Setup wizard 248
- DB2 pureScale Feature (*continued*)
 - installing (*continued*)
 - db2_install command 267
 - methods 246
 - overview 143
 - preinstallation tasks on AIX 161
 - preinstallation tasks on Linux 195
 - preinstallation tasks overview 161
 - prerequisites for AIX 161, 174
 - prerequisites for Linux 195, 206
 - verifying installation 273
 - post-conversion tasks
 - database configuration settings 452
 - overview 451
 - pre-conversion tasks
 - converting unsupported table spaces 445
 - overview 443
 - setting up a GPFS file system 444
 - verifying databases 446
 - preferred primary cluster caching facilities 248
 - preferred secondary cluster caching facilities 248
 - removing 313
 - taking over a user-managed GPFS cluster 274
 - uninstalling 313
- DB2 pureScale instances
 - converting 449
 - creating 276
 - extending 280
 - updating 285
- DB2 servers
 - installing
 - Linux 13
 - UNIX 13
 - Windows 95
 - installing walkthrough 103
 - overview 11
 - partitioned
 - Windows 323
 - pre-conversion tasks
 - backing up databases 447, 452
- DB2 services
 - Windows 51
- DB2 Setup wizard
 - creating response files 373
 - DB2 pureScale Feature
 - overview 246
 - procedure 248
 - installing
 - DB2 servers (Linux)DB2 servers (UNIX) 99
 - language identifiers 82
- db2_deinstall command
 - dropping cluster caching facility 313
 - dropping member 313
 - removing DB2 products 494
- db2_install command
 - DB2 pureScale Feature 246, 267
- DB2ADMNS group
 - adding users 123
- db2cluster command
 - verifying configuration of file system cluster 155
- db2icrt command
 - creating DB2 pureScale instances
 - procedure 276
 - user-managed file systems 155
 - creating instances 37
 - installing DB2 pureScale Feature 267

- db2instance command
 - verifying DB2 pureScale installation and setup 273
- db2isetup command
 - creating DB2 pureScale instance 276
 - creating response files 373
 - DB2 pureScale environment 280
- db2iupdt command
 - DB2 pureScale environments
 - adding cluster caching facility 280
 - adding members 280
 - converting to DB2 pureScale instance 449
 - dropping member 313
 - installing DB2 pureScale Feature 267
 - updating instance to higher level within release 285
 - updating to DB2 pureScale instance 281
 - user-managed file systems 155
 - dropping cluster caching facility 313
 - updating 32-bit instances to 64-bit instances 124
- db2licm command
 - registering licenses 276
- db2ls command
 - listing installed products and features 139
- db2nodes.cfg file
 - format 349
 - netname field 323
 - updating 348
- db2osconf command
 - determining kernel configuration parameter values 43
- db2prereqcheck command
 - checking installation prerequisites 3
- db2rfe command
 - enabling root features 19, 21
- db2setup command
 - DB2 pureScale Feature 246, 248
 - language setting 82
- DB2USERS user group
 - adding users 123
- db2val command
 - validating DB2 copy 124, 141
- default DB2 copy
 - changing after installation 125
- default IBM database client interface copy
 - changing after installation 125
- default language setting
 - Windows 84
- directories
 - shared access 382
- directory schema
 - extending
 - Windows 55
- directory structures
 - Linux 78
 - root installations compared to non-root installations 20
 - Windows 72
- disk space
 - requirements 31
- distributed installation
 - DB2 products
 - SCCM 384
- doce_deinstall command
 - removing products 494
- documentation
 - Information Center
 - installing 427
 - overview 551
 - PDF files 551
 - printed 551

- documentation (*continued*)
 - terms and conditions of use 558
- DVDs
 - mounting
 - AIX 91
 - HP-UX 92
 - Linux 93
 - Solaris Operating Environment 93

E

- errors
 - response file error codes
 - Linux 378
 - UNIX 378
 - Windows 385
- exports
 - profiles 413

F

- fast communication manager
 - See FCM 34
- FCM
 - communications between database partition servers 355
 - overview
 - Linux 34
 - UNIX 34
 - port numbers 355
- fcm_num_buffers configuration parameter
 - overview 34
- fcm_num_channels configuration parameter
 - overview 34
- fenced users 35
- file system clusters
 - verifying configuration 155
- file systems
 - creating for partitioned database system
 - Linux 332
 - Solaris Operating Environment 333
 - creating for partitioned DB2 server 328
- Firefox browser profile deletion 121
- First Steps
 - deleting Firefox profile 121
 - overview 122
 - product updates 138
- fix packs
 - downloading 461
 - installing
 - Linux 469
 - methods 467
 - Microsoft Cluster Server environment 477
 - multiple products (Windows) 475
 - new products (Linux and UNIX) 474
 - overview 282, 455
 - prerequisites 459
 - response file (Windows) 476
 - single product (Windows) 475
 - UNIX 469
 - updating DB2 pureScale instance 469
 - updating products (Linux and UNIX) 473
 - Windows 475
 - post-installation
 - binding manually 484
 - tasks (Linux) 482
 - tasks (UNIX) 482

- fix packs (*continued*)
 - post-installation (*continued*)
 - tasks (Windows) 481
 - prerequisites 460, 463, 464
 - uncompressing 462
 - uninstalling 487, 513
 - universal versus product-specific 468

G

GDPCs

- configuring 293
- creating databases 309
- HADR
 - configuring 309
- installing 294
- prerequisites 242
- setting up GPFS replication 302

globalization

- interface language
 - DB2 Setup wizard 81

GPFS

- pre-conversion tasks for DB2 pureScale environments 444
- replication 302

groups

- creating IDs 35
- Linux 32
- overview 32
- partitioned database environments 336
- UNIX 32

H

hardware

- requirements
 - DB2 database products 1
 - DB2 server products (AIX) 56
 - DB2 server products (HP-UX) 58
 - DB2 server products (Linux) 59
 - DB2 server products (Solaris Operating Environment) 66
 - IBM data server clients (AIX) 56
 - IBM data server clients (HP-UX) 58
 - IBM data server clients (Linux) 59
 - IBM data server clients (Solaris Operating Environment) 66
 - Windows 68

help

- SQL statements 554

home file system

- AIX 328
- HP-UX 330
- overview 328

HP-UX

- DB2 home file system 330
- installing
 - DB2 servers 13, 58
 - IBM data server clients 58
- kernel configuration parameters
 - modifying 45
 - recommended values 43
- mounting media 92
- Network File System (NFS) 335
- partitioned database systems 330
- user creation 338

I

I/O

- fast I/O fencing 151

IBM data server clients

- connectivity using multiple DB2 copies 135
- overview 11

IBM Data Server Driver for JDBC and SQLJ

- levels for DB2 versions 27

IBM Data Studio

- DB2 integration 543
- installation scenario 547
- installing with DB2 Setup wizard
 - Linux 545
 - Windows 543

IBM database client interface copies

- changing default copy after installation 125
- uninstalling 501

IBM Secure Shell Server for Windows

- installing 139

IBM Tivoli System Automation for Multiplatforms (SA MP)

- install log 535
- installing
 - DB2 installer 515, 520
 - overview 519
 - uninstallSAM script 521
- license terms 537
- overview 517
- system requirements 539
- uninstall log 535
- uninstalling
 - DB2 installer 527
 - overview 527
 - uninstallSAM script 528
- upgrading
 - DB2 installer 524
 - DB2 pureScale environment 472
 - overview 523
 - uninstallSAM script 525

IBM Tivoli System Automation for Multiplatforms (SA MP)

automated failover scripts

- installing
 - DB2 installer 531
 - manually 532
 - overview 531
- uninstalling
 - DB2 installer 531
 - manually 532
 - overview 531

upgrading

- DB2 installer 531
- manually 532
- overview 531

imports

- profiles 413

installation

- add-on products 137
- database partition servers
 - response files (Linux) 360, 378
 - response files (overview) 359
 - response files (UNIX) 360, 378
 - response files (Windows) 359
- database servers 31, 32
- DB2 database user accounts (Windows) 48
- DB2 Enterprise Server Edition 323, 328
- DB2 Information Center
 - DB2 Setup wizard (Linux) 421
 - DB2 Setup wizard (Windows) 425

- installation (*continued*)
 - DB2 Information Center (*continued*)
 - options 427
 - DB2 products
 - as non-root user 17
 - methods that require manual configuration 41
 - DB2 pureScale Feature
 - AIX user requirements 170
 - DB2 Setup wizard 248
 - first steps after installing 273
 - methods 246
 - overview 143
 - prerequisites 161, 195
 - verifying 273
- fix packs
 - bind files 484
 - DB2 pureScale instance 469
 - methods 467
 - Microsoft Cluster Server environment 477
 - multiple database products 475
 - new DB2 database products 474
 - obtaining 461
 - post-installation tasks 481, 482
 - pre-installation tasks 459
 - prerequisites 460, 463, 464
 - response file 476
 - single database product 475
 - types 468
 - uncompressing 462
 - updating database products 473
- images
 - embedding in application 411
- interactive 415
- kernel parameters 42
- listing DB2 database products 139
- mounting product media 91
- multiple DB2 copies 90
- partitioned database environments
 - verifying 343
- post-installation tasks
 - Linux 139
 - overview 121
 - UNIX 139
 - Windows 123
- prerequisites
 - DB2 pureScale 161
 - DB2 pureScale environments 195
 - Linux 32
 - overview 32
 - UNIX 32
 - Windows 48
- requirements
 - AIX 39, 56
 - HP-UX 58
 - Linux 59
 - overview 1
 - Solaris Operating Environment 66
- response files
 - error codes 378, 385
 - making the DB2 installation files available 381
 - overview 365
 - types 363
- security 34
- silent
 - Linux 375
 - UNIX 375
 - Windows 381

- installation (*continued*)
 - updating AIX environment settings 325
 - verifying
 - using CLP 119
 - Windows elevated privileges 53
 - zSeries running Linux
 - DB2 35
- installing
 - requirements
 - Windows 68
- instances
 - converting
 - between DB2 database products (Linux) 437
 - between DB2 database products (UNIX) 437
 - between DB2 database products (Windows) 435
 - to DB2 pureScale environment 449
 - creating
 - using db2icrt command 37
 - dropping 493
 - removing 493, 498
 - stopping
 - Linux 492
 - non-root 497
 - UNIX 492
- interface languages
 - changing
 - UNIX 83
 - Windows 84
 - DB2 Setup wizard 81
 - overview 82

J

- Java
 - DB2 product support 27
- JDBC
 - drivers
 - details 27

K

- kernel configuration parameters
 - HP-UX
 - db2osconf command 43
 - modifying 45
 - recommended 43
 - Linux
 - modifying 46, 62, 232
 - requirements 43, 64
 - Solaris 47

L

- LANG environment variable
 - setting 82, 83
- languages
 - DB2 interface 82, 84
 - DB2 Setup wizard for language identifiers 82
- licenses
 - applying 122
 - registering
 - db2licm command 276
- Linux
 - changing DB2 database product editions 437
 - default port ranges 355

- Linux (*continued*)
 - installing
 - DB2 databases on zSeries 35
 - DB2 pureScale environments 195
 - DB2 servers 13, 59, 99
 - IBM data server clients 59
 - IBM Data Studio with DB2 Setup wizard 545
 - response file 375
 - kernel parameters
 - modifying 46, 62, 232
 - requirements 43, 64
 - libraries
 - libaio.so.1 59
 - libstdc so.5 59
 - listing DB2 database products 139
 - mounting
 - CDs 93
 - DVDs 93
 - NFS verification 335
 - partitioned database system file systems 332
 - removing
 - DB2 (root) 491
 - required users 203, 339
 - uninstalling DB2
 - non-root 498
 - root 491, 493
 - verifying partitioned database server installations 343
- locales
 - DB2 interface languages 82

M

- members
 - adding 280
- memory
 - requirements
 - overview 31
- monitoring
 - DB2 pureScale environments
 - shared file systems 291
 - shared file systems in DB2 pureScale environments 291
- mounting CDs or DVDs
 - AIX 91
 - HP-UX 92
 - Linux 93
 - Solaris Operating Environment 93

N

- national language support (NLS)
 - displaying DB2 Setup wizard 82
- Network File System (NFS)
 - installing DB2 41
 - verifying operation 335
- Network Time Protocol
 - see NTP 70
- node configuration files
 - format 349
 - updating 348
- non-root installations
 - differences 20
 - directory structure 20
 - enabling root-based features 19
 - installing 17
 - limitations 21
 - overview 17

- non-root installations (*continued*)
 - ulimit settings 38
 - uninstalling 498
- non-root instances
 - dropping 498
 - removing 498
- notices 561
- notification lists
 - setting 124
- NTP
 - setting up
 - configuring hosts as NTP clients 70, 270
 - NTP server 71, 269
 - overview 70

O

- openSSH 91
- OpenSSH 41, 172, 204
- operating system requirements
 - Windows 68
- operating systems
 - installation requirements
 - AIX 56
 - HP-UX 58
 - Linux 59
 - Solaris Operating Environment 66
 - prerequisites 1

P

- paging space
 - requirements 31
- Partitioned database environment
 - install 317
- partitioned database environments
 - installing 323
- partitioned databases
 - groups 336
 - installation verification
 - Linux 343
 - overview 343
 - UNIX 343
 - Windows 344
 - post-installation tasks
 - overview 347
 - setting up 319
 - users 336
- passwordless SSH 41, 172, 204
- passwords
 - rules 85
- performance
 - DB2 cluster services
 - fencing 151
- planning
 - shared storage 151
- port number ranges
 - enabling communication
 - Linux 355
 - UNIX 355
 - verifying availability
 - Linux 327
 - UNIX 327
- post-conversion tasks
 - DB2 pureScale environments
 - overview 451

- post-conversion tasks *(continued)*
 - DB2 pureScale Feature
 - database configuration settings 452
- post-installation tasks
 - Linux 139
 - Main menu entries 141
 - overview 121
 - partitioned database environments 347
 - UNIX 139
 - Windows 123
- pre-conversion tasks
 - DB2 pureScale environments 443
 - DB2 pureScale Feature 444
 - converting unsupported table spaces 445
 - verifying databases are ready for conversion 446
 - DB2 servers
 - backing up databases 447, 452
- preferred cluster caching facilities
 - specifying 248
- preinstallation cheat sheet 190, 234
- preinstallation checklist 190, 234
- preinstallation tasks
 - DB2 pureScale Feature
 - overview 161, 195
- problem determination
 - information available 558
 - tutorials 558
- processes
 - stopping
 - interactive installation 415
 - response file installation 417
- profiles
 - exporting 413
 - importing 413
- public key-based authentication 41, 172, 204
- PVIDs 188

R

- registry
 - verifying access 343
- remote commands
 - enabling (UNIX) 357
- removing
 - DB2 (root)
 - Linux 491
 - UNIX 491
 - DB2 pureScale Feature 313
- response files
 - creating
 - DB2 Setup wizard 373
 - UNIX 376
 - Windows 382
 - DB2 pureScale
 - installation overview 246
 - DB2 pureScale Feature
 - special considerations 369
 - exporting configuration profile 413
 - generator
 - overview 387
 - importing configuration profile 413
 - installation
 - batch files 385
 - database partition servers 359, 360, 378
 - DB2 pureScale 266, 376
 - killing DB2 processes 417
 - Linux 375

- response files *(continued)*
 - installation *(continued)*
 - making DB2 files available for 381
 - types 363
 - UNIX 375
 - Windows 381
 - installation error codes 378, 385
 - keywords
 - details 393
 - location 365, 367, 373
 - overview 246, 365, 367
 - running setup 383
 - sample 389, 391
 - uninstallation
 - Linux 379, 509
 - UNIX 379, 509
 - Windows 386, 511
- rhosts file
 - creating 357
- root installations
 - comparison with non-root installations 20
 - directory structure 20
- root instances
 - dropping 493
 - removing 493
- root-based features
 - non-root installations 19
- rsh 357

S

- samples
 - response files 389, 391
- SCCM
 - distributed installation
 - DB2 products 384
- SDKs
 - product levels 27
- secondary cluster caching facilities
 - specifying 248
- security
 - user groups 123
- setting up
 - 10GE on Linux 221
 - InfiniBand on AIX 178
 - InfiniBand on Linux 208
- setting up InfiniBand 208
- setting up uDAPL 178, 208, 221
- shared file systems
 - creating 289
 - deleting 290
 - monitoring 291
- silent installation
 - keywords 393
 - Linux 375
 - overview 365
 - UNIX 375
 - Windows 381
- silent uninstallation
 - Linux 379, 509
 - overview 367
 - UNIX 379, 509
 - Windows 386, 511
- software requirements
 - AIX 56
 - HP-UX 58
 - Linux 59

- software requirements (*continued*)
 - overview 1
 - Solaris Operating Environment 66
 - Windows 68
- Solaris operating systems
 - creating file systems 333
 - creating required users 340
 - installation requirements
 - DB2 servers 66
 - IBM data server clients 66
 - installing
 - DB2 servers 13
 - modifying kernel parameters 47
 - mounting CDs or DVDs 93
 - verifying that NFS is running 335
- SQL statements
 - help
 - displaying 554
- ssh 357
- starting
 - Information Center 429
- stopping
 - DB2 administration server 491
 - Information Center 429
 - root instances 492
- storage
 - planning shared storage 151
- swap space
 - See* paging space
- system requirements
 - DB2 Database for Linux, UNIX, and Windows 1
- Systems Center Configuration Manager
 - distributed installation
 - DB2 products 384

T

- table space conversion
 - pre-conversion tasks for DB2 pureScale environments 445
- terms and conditions
 - publications 558
- troubleshooting
 - online information 558
 - tutorials 558
- tutorials
 - list 558
 - problem determination 558
 - pureXML 558
 - troubleshooting 558

U

- uDAPL communication 178, 208, 221
- ulimit settings 38
- uninstallation
 - DAS 492
 - DB2 copies 501
 - DB2 database products
 - Windows 489, 499
 - DB2 pureScale Feature 313
 - db2_deinstall command 494
 - doce_deinstall command 494
 - fix packs 487, 513
 - IBM database client interface copies 501
 - non-root installations 497, 498
 - root installations 491

- uninstallation (*continued*)
 - using response files 367, 379, 386, 509, 511
- uninstalling
 - root installations 491
- UNIX
 - changing DB2 database product editions 437
 - changing DB2 interface language 83
 - DB2 groups 32
 - DB2 users 32
 - default port ranges 355
 - installing
 - DB2 Setup wizard 99
 - listing DB2 database products 139
 - partitioned database server installation verification 343
 - remote command execution 357
 - removing
 - DB2 (root) 491
 - DB2 non-root installations 497
 - DB2 non-root instances 498
 - DB2 products 494
 - DB2 root instances 493
 - response file creation 376
 - response file installation 375
 - uninstalling DB2
 - root 491
 - updating node configuration file 348
- updates
 - 32-bit DB2 instances 124
 - 64-bit DB2 instances 124
 - checking for 138
 - DB2 Information Center 431, 555, 556
 - DB2 pureScale instances 285
 - db2nodes.cfg (UNIX) 348
 - instance to DB2 pureScale 281
 - node configuration file 348
 - rhhosts file 357
- user accounts
 - DB2 Administration Server (Windows) 48
 - db2fmp process (Windows) 52
 - instance user (Windows) 48
 - required for installation (Windows) 48
- user groups
 - DB2ADMNS 123
 - DB2USERS 123
 - security 123
- user IDs
 - creating 35
- user limit requirements 38
- user privileges
 - Windows 55
- user-managed GPFS cluster
 - manually allowing the DB2 installer to take over 274
- users
 - creating required
 - AIX 336
 - HP-UX 338
 - Linux 203, 339
 - Solaris Operating System 340
 - overview 32
 - partitioned database environments 336
- utilities
 - response file generator 387

V

- validation
 - DB2 copies 124, 141

- verification
 - port range availability
 - Linux 327
 - UNIX 327
 - registry access 343

W

Windows

- DB2 services 51
- default language setting 84
- elevated privileges setup 53
- granting user rights 55
- installation verification
 - partitioned database environments 344
- installing
 - DB2 servers (procedure) 15
 - DB2 servers (with DB2 Setup wizard) 95
 - IBM Data Studio with DB2 Setup wizard 543
 - response file method 383
- response files
 - editing 382
 - installing using 381
- shared directory access setup 382
- uninstalling DB2 499
- user accounts
 - DB2 server product installation 48
 - db2fmp process 52

Windows operating systems

- installing
 - DB2 servers (requirements) 68
 - IBM data server clients (requirements) 68

workload balancing 159

workload partitions

- installing 39

Z

zSeries

- installing DB2 Database for Linux 35



Printed in USA

GC27-3884-01



Spine information:

IBM DB2 10.1 for Linux, UNIX, and Windows

Installing DB2 Servers

